

E. E. Holmes and E. J. Ward

# Analysis of multivariate time-series using the MARSS package

version 2.1

May 26, 2011

Mathematical Biology Program  
Northwest Fisheries Science Center, Seattle, WA

Holmes, E. E. and E. J. Ward. 2010. Analysis of multivariate time-series using the MARSS package. NOAA Fisheries, Northwest Fisheries Science Center, 2725 Montlake Blvd E., Seattle, WA 98112. Contacts [eli.holmes@noaa.gov](mailto:eli.holmes@noaa.gov) and [eric.ward@noaa.gov](mailto:eric.ward@noaa.gov)

## Preface

The initial motivation for our work with MARSS models was a collaboration with Rich Hinrichsen. Rich developed a framework for analysis of multi-site population count data using MARSS models and bootstrap AICb (Hinrichsen and Holmes, 2009). Our work (EEH and EJW) extended Rich's framework, made it more general, and led to the development of a parametric bootstrap AICb for MARSS models, which allows one to do model-selection using datasets with missing values (Ward et al., 2009; Holmes and Ward, 2010). Later, we developed additional algorithms for simulation and confidence intervals. Discussions with Mark Scheuerell led to an extensive revision of the EM algorithm and to the development of a general EM algorithm for constrained MARSS models (Holmes, 2010). Discussions with Mark also led to a complete rewrite of the model specification so that the package could be used for MARSS models in general—rather than simply the form of MARSS model used in our (EEH and EJW) applications. Many collaborators have helped test the package; we thank especially Yasmin Lucero, Mark Scheuerell, Kevin See, and Brice Semmens. Development of the code into a R package would not have been possible without Kellie Wills, who wrote the much of the code outside of the algorithm functions.

The case studies used in this manual were developed for workshops on analysis of multivariate time-series data given at the Ecological Society meetings since 2005 and taught by us (EEH and EJW) along with Yasmin Lucero, Stephanie Hampton, Brice Semmens, and Mark Scheuerell. The case study on extinction estimation and trend estimation was initially developed by Brice Semmens and later extended by us for this manual. The algorithm behind the TMU figure in Chapter 8 was developed during a collaboration with Steve Ellner (Ellner and Holmes, 2008).

EEH and EJW are research scientists at the Northwest Fisheries Science Center in the Mathematical Biology program. This work was conducted as part of our jobs at the Northwest Fisheries Science Center, a research center for NOAA Fisheries which is a US federal government agency. A CAMEO grant from NOAA Fisheries supported Kellie Wills. During the initial stages of this work, EJW was supported on a post-doctoral fellowship from the National Research Council.

You are welcome to use the code and adapt it with attribution. It may not be used in any commercial applications. Links to more code and publications on MARSS applications can be found by following the links at EEH's website <http://faculty.washington.edu/eeholmes>. Links to our papers that use these methods can also be found at the same website.



---

## Contents

<b>1</b>	<b>The MARSS package</b>	<b>1</b>
1.1	What does the MARSS package do?	2
1.2	How to get started (quickly)	2
1.3	Important notes about the algorithms	3
1.4	Troubleshooting	5
1.5	Other related packages	6
<b>2</b>	<b>Overview of the package functions</b>	<b>9</b>
2.1	The <code>MARSS()</code> function	9
2.2	Core functions for fitting a MARSS model	10
2.3	Functions for a fitted <code>marssMLE</code> object	10
2.4	Functions for <code>marssm</code> objects	11
<b>3</b>	<b>The <code>MARSS()</code> function</b>	<b>13</b>
3.1	Process equation constraint options	14
3.2	Observation equation constraints	18
<b>4</b>	<b>Model specification in the core functions</b>	<b>21</b>
4.1	The fixed and free components of the model parameters	21
4.2	Limits on the model forms that can be fit	23
<b>5</b>	<b>Algorithms used in the MARSS package</b>	<b>25</b>
5.1	Kalman filter and smoother	25
5.2	The exact likelihood	26
5.3	Maximum-likelihood parameter estimation	27
5.4	Parametric and innovations bootstrapping	28
5.5	Simulation and forecasting	29
5.6	Model selection	29

<b>6</b>	<b>Examples</b>	31
6.1	Fitting different MARSS models to a dataset	31
6.2	Printing and summarizing models and model fits	38
6.3	Confidence intervals on a fitted model	39
6.4	Vectors of just the estimated parameters	40
6.5	Degenerate variance estimates	41
6.6	Bootstrap parameter estimates	43
6.7	Data simulation	43
6.8	Bootstrap AIC	46
<b>7</b>	<b>Case study instructions</b>	47
<b>8</b>	<b>Case Study 1: Count-based Population Viability Analysis using corrupted data</b>	49
8.1	The Problem	49
8.2	Simulated data with process and observation error	50
8.3	Maximum-likelihood parameter estimation	53
8.4	Probability of hitting a threshold $\Pi(x_d, t_e)$	58
8.5	Certain and uncertain regions	62
8.6	More risk metrics and some real data	64
8.7	Confidence intervals	66
8.8	Comments	67
<b>9</b>	<b>Case study 2: Combining multi-site and subpopulation data to estimate regional population dynamics</b>	69
9.1	The problem	69
9.2	Analysis assuming a single total Puget Sound population	71
9.3	Changing the assumption about the observation variances	76
9.4	Analysis assuming north and south subpopulations	79
9.5	Using MARSS() to fit other population structures	83
9.6	Discussion	85
<b>10</b>	<b>Case Study 3: Using MARSS models to identify spatial population structure and covariance</b>	89
10.1	The problem	89
10.2	How many distinct subpopulations?	90
10.3	Is Hood Canal separate?	93
<b>11</b>	<b>Case Study 4: Dynamic factor analysis (DFA) using MARSS</b>	97
11.1	Dynamic factor analysis	97
11.2	The data	99
11.3	Setting up the model	100
11.4	Fitting up the model	103
11.5	Using model selection to determine the number of trends	105
11.6	Using a varimax rotation to determine the trend loadings	109

11.7 Discussion .....	109
<b>12 Case Study 5: Using state-space models to analyze noisy animal tracking data</b> .....	111
12.1 A simple random walk model of animal movement .....	111
12.2 The problem .....	112
12.3 Estimate locations from bad tag data .....	112
12.4 Comparing turtle tracks to proposed fishing areas .....	116
12.5 Using specialized packages to analyze tag data .....	117
<b>A Textbooks and articles that use MARSS modeling for population modeling</b> .....	119
<b>B Package MARSS: Object structures</b> .....	123
B.1 Model objects: class marssm .....	123
B.2 Wrapper objects: class popWrap .....	124
B.3 ML estimation objects: class marssMLE .....	125
<b>C Package MARSS: The top-level MARSS functions and the base functions</b> .....	127
<b>References</b> .....	129
<b>Index</b> .....	133





## The MARSS package

MARSS stands for Multivariate Auto-Regressive(1) State-Space. The MARSS package is designed for linear MARSS models with Gaussian errors. This class of model is extremely important in the study of linear stochastic dynamical systems, and these models are important in many different fields, including economics, engineering, genetics, physics and ecology. Appendix A gives a selection of textbooks on MARSS models.

A MARSS model, with Gaussian errors, takes the form:

$$\mathbf{x}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{u} + \mathbf{w}_t, \text{ where } \mathbf{w}_t \sim \text{MVN}(0, \mathbf{Q}) \quad (1.1a)$$

$$\mathbf{y}_t = \mathbf{Z}\mathbf{x}_t + \mathbf{a} + \mathbf{v}_t, \text{ where } \mathbf{v}_t \sim \text{MVN}(0, \mathbf{R}) \quad (1.1b)$$

$$\mathbf{x}_1 \sim \text{MVN}(\boldsymbol{\pi}, \mathbf{V}_1) \quad (1.1c)$$

The model includes random variables, parameters and data:

$\mathbf{x}_t$  is a  $m \times 1$  column vector of the hidden states at time  $t$ . It is a realization of the random variable  $\mathbf{X}_t$ .

$\mathbf{w}_t$  is a  $m \times 1$  column vector of the process errors at time  $t$ . It is a realization from a multivariate normal random variable with mean 0 and  $\Sigma = \mathbf{Q}$ .

$\mathbf{y}_t$  is a  $n \times 1$  column vector of the observed data at time  $t$ .

$\mathbf{v}_t$  is a  $n \times 1$  column vector of the non-process errors at time  $t$ . It is a realization from a multivariate normal random variable with mean 0 and  $\Sigma = \mathbf{R}$ .

$\mathbf{B}$  is a parameter and is a  $m \times m$  matrix.

$\mathbf{u}$  is a parameter and is a  $m \times 1$  column vector.

$\mathbf{Q}$  is a parameter and is a  $m \times m$  variance-covariance matrix.

$\mathbf{Z}$  is a parameter and is a  $n \times m$  matrix.

$\mathbf{a}$  is a parameter and is a  $n \times 1$  column vector.

$\mathbf{R}$  is a parameter and is a  $n \times n$  variance-covariance matrix.

$\boldsymbol{\pi}$  is either a parameter or a fixed prior. It is a  $m \times 1$  matrix.

$\mathbf{V}_1$  is a fixed value. It is a  $m \times m$  variance-covariance matrix.

The meaning of the parameters in the MARSS models depends on the application for which the MARSS model is being used. In the case studies, we

show examples of MARSS models used to analyze population count data and animal tracking data, and Appendix A gives a selection of papers from the ecological literature. However, the MARSS package is not specific to population modeling applications. The functions in the MARSS package are generic functions for fitting MARSS models have forms like:

## 1.1 What does the MARSS package do?

The MARSS package is designed to fit unconstrained and constrained MARSS models. A constrained MARSS model is one in which some of the parameters are constrained in the sense that they have fixed, free and/or shared values. For example, let  $\mathbf{M}$  and  $\mathbf{m}$  be arbitrary matrix and column vector parameters. The MARSS package allows one to specify and fit models where  $\mathbf{M}$  and  $\mathbf{m}$  have shared values and fixed values. For example,

$$\mathbf{M} = \begin{bmatrix} a & 0.9 & c \\ -1.2 & a & 0 \\ 0 & c & b \end{bmatrix} \text{ and } \mathbf{m} = \begin{bmatrix} d \\ d \\ e \\ 2.2 \end{bmatrix}$$

Version 1.x of the MARSS package fits models via maximum-likelihood using a Kalman-EM algorithm<sup>1</sup>. The Kalman-EM algorithm is used because it gives robust estimation for datasets replete with missing values and for models with various constraints. The MARSS package also supplies functions for bootstrap and approximate confidence intervals, parametric and non-parametric bootstrapping, model selection (AIC and bootstrap AIC), simulation, and bootstrap bias correction. Version 1.0 does not allow  $\mathbf{B}$  or  $\mathbf{Z}$  to be estimated. Version 2.0 is currently being tested and it will allow  $\mathbf{B}$  and  $\mathbf{Z}$  estimation along with less constrained forms of  $\mathbf{Q}$  and  $\mathbf{R}$ .

## 1.2 How to get started (quickly)

Install the MARSS package and then type `library(MARSS)` at the command line to load the package. Read the first 2-4 pages of Chapter 3 then read through a couple of the examples in Chapter 6. Get your data into a matrix (not dataframe) with time going across the columns and any non-*yy* columns (like year) removed. Replace any missing time steps with a missing value holder (like NA or -99). Write your model down on paper and identify which

---

<sup>1</sup> The package can also fit models via quasi-Newton methods based on *R*'s `optim` function. This can be especially useful for finishing off a Kalman-EM estimate when the data to parameter ratio is high. However, when the ratio of data to parameters is low (as in many ecological applications), the quasi-Newton algorithm tends to be fragile and sensitive to initial conditions.

parameters correspond to  $\mathbf{B}$ ,  $\mathbf{u}$ ,  $\mathbf{Q}$ ,  $\mathbf{Z}$ ,  $\mathbf{a}$ , and  $\mathbf{R}$  in the MARSS model (Equation 1.1). Call the `MARSS()` function (Chapter 3) using your data and using the `constraint` argument to specify the form of the parameters (Chapter 3 shows the options for the forms).

### 1.3 Important notes about the algorithms

MARSS provides maximum-likelihood via an EM algorithm using the Kalman filter/smoother. All code is in native *R*. Thus the model fitting is slow (relatively)<sup>2</sup>. EM algorithms will quickly get in the vicinity of the maximum likelihood, but the final approach to the maximum is generally slow relative to quasi-Newton methods. On the flip side, EM algorithms are quite robust to initial conditions choices. The MARSS package allows one to use the BFGS method in `optim()` to fit MARSS models. The *DLM R* package (search for it on CRAN) also provides estimation of MARSS models via quasi-Newton methods as does the *SsfPack* package (just do an internet search for that package).

Restricted maximum-likelihood algorithms are also available for AR-1 state-space models, both univariate (Staples et al., 2004) and multivariate (Hinrichsen, 2009). REML can give parameter estimates with lower variance than plain maximum-likelihood algorithms. However, the algorithms for REML when there are missing values are not currently available (although that will probably change in the near future). Another maximum-likelihood method is data-cloning which adapts MCMC algorithms used in Bayesian analysis for maximum-likelihood estimation (Lele et al., 2007).

Data with cycles, from say internal dynamical interactions, are difficult to analyze, and both REML and Kalman-EM approaches will give poor estimates for this type of data. The slope method (Holmes, 2001), is more ad-hoc but is relatively robust to those problems. Holmes et al. (2007) used the slope method in a large study of data from endangered and threatened species; Ellner and Holmes (2008) showed that the slope estimates are close to the theoretical minimum uncertainty. However estimates using the slope method are not easily extended to multi-variate data and it is not a true maximum-likelihood method.

Missing values are seamlessly accommodated with the MARSS package. Simply specify the way missing values are denoted in the data set (default is

---

<sup>2</sup> In addition, the Kalman filter/smoothing algorithm in MARSS is implemented exactly as you see it in Shumway and Stoffer (2006, p. 331-335) (based on the original smoother presented in Rauch (1963)). Table 2 in Koopman (1993) indicates that this algorithm is 20-40 times less efficient than more optimal formulations; see also (Kohn and Ansley, 1989; Koopman et al., 1999). MARSS 2.2 uses the Kalman filter/smoothing from the *KFAS* package when possible to speed up estimation; this filter/smoothing uses a more efficient algorithm than that in the `MARSSkf()` function and is 40-100 times faster.

`miss.value=-99`). The likelihood computations are exact and will deal appropriately with missing values. The presence of missing values, however, limits the  $\mathbf{R}$  matrix to being a diagonal matrix (if estimated) and the  $\mathbf{Z}$  matrix to being fixed. In addition, no innovations<sup>3</sup> bootstrapping can be done if there are missing values. Instead parametric bootstrapping must be used.

You should be aware that maximum-likelihood estimates of variance in MARSS models are fundamentally biased, regardless of the algorithm used. This bias is more severe when one or the other of  $\mathbf{R}$  or  $\mathbf{Q}$  is very small, and the bias does not go to zero as sample size goes to infinity. The bias arises because variance is constrained to be positive. Thus if  $\mathbf{R}$  or  $\mathbf{Q}$  is essentially zero, the mean estimate will not be zero and thus the estimate will be biased high while the corresponding bias of the other variance will be biased low. You can generate unbiased variance estimates using a bootstrap estimate of the bias. The function `MARSSparamCIs()` will do this. However be aware that adding an *estimated* bias to a parameter estimate will lead to an increase in the variance of your parameter estimate. The amount of variance added will depend on sample size.

You should also be aware that mis-specification of the prior on the initial states can have catastrophic affects on your parameter estimates if your prior conflicts with the distribution of the initial states implied by the MARSS model. These affects can be very difficult to detect because the model will appear to well-fitted. Unless you have a good idea of what the parameters should be, you might know that your prior conflicts. The most common problems, we have found are the following. 1) The correlation structure in the prior (whether diffuse or not) does match the correlation structure implied by the MARSS model. For example, you specify a diagonal  $V_0$  (independent states), but the implied distribution has correlations. 2) The correlation structure in the prior does not match the structure implied by constraints on  $\mathbf{x}_0$ . For example, you specify that all values in  $\mathbf{x}_0$  are shared, yet you specify that  $V_0$  is diagonal (independent). 3) You specify that  $V_0=0$  and estimate  $\mathbf{x}_0^0$  for your prior, but your model does not "run backwards" robustly so you cannot estimate  $\mathbf{x}_0^0$ . For example, you have a model with a non-diagonal  $\mathbf{B}$  matrix, this model needs the  $\mathbf{x}$  to be constrained by data, but  $\mathbf{x}_0^0$  has no  $\mathbf{y}_0$  to constrain it. You might get around this by specifying your prior as  $\mathbf{x}_1^0$ . You can get around these problems sometimes by using a diffuse prior (using `method="BFGS"` and `control$diffuse=TRUE`) but you do not know the correlation structure of  $\mathbf{x}_1^0$  then your diffuse prior will conflict with the model. The other way to get around the problem is to set  $V_0=0$  and estimate  $\mathbf{x}_1^0$  or  $\mathbf{x}_0^0$  as the prior. Whether to use  $\mathbf{x}_1^0$  or  $\mathbf{x}_0^0$  depends on the model.

---

<sup>3</sup> referring to the non-parametric bootstrap used in the package

## 1.4 Troubleshooting

There are two numerical errors and warnings that you may see when fitting MARSS models: ill-conditioning and degeneracy. The Kalman and EM algorithms need inverses of matrices. If those matrices become ill-conditioned, for example all elements are close to the same value, then the algorithm becomes unstable. MARSS will print warning messages if the algorithm is becoming unstable and you can set `control$trace=1`, to see details of where the algorithm is becoming unstable. Whenever possible, you should avoid using shared  $\pi$  values in your model<sup>4</sup>. The way our algorithm deals with  $\mathbf{V}_1$  tends to make this case unstable, especially if  $\mathbf{R}$  is not diagonal. In general, estimation of a non-diagonal  $\mathbf{R}$  is more difficult, more prone to ill-conditioning, and more data-hungry.

The second numerical error you may see is a degeneracy warning. This means that one of the elements on the diagonal of your  $\mathbf{Q}$  or  $\mathbf{R}$  matrix are going to zero (are degenerate). It will take the EM algorithm forever to get to zero. BFGS will have the same problem, although it will often get a bit closer to the degenerate solution. If you are using `method="kem"`, MARSS will warn you if it looks like the solution is degenerate. If you use `control=list(allow.degen=TRUE)`, the EM algorithm will attempt to set the degenerate variances to zero (instead of trying to get to zero using an infinite number of iterations). If you still get the degenerate warnings, look at your data for anything obvious, like having only one data point in your time series or some really extreme outlier. Then try running the EM algorithm longer using `control=list(maxit=something big)`.

The algorithms in the MARSS 1.0 package are designed for cases where the  $\mathbf{Q}$  and  $\mathbf{R}$  diagonals are all non-minuscule. For example, the EM update equation for  $\mathbf{U}$  will grind to a halt (not update  $\mathbf{U}$ ) if  $\mathbf{Q}$  is tiny (like 1E-7). Conversely, the BFGS equations are likely to miss the maximum-likelihood when  $\mathbf{R}$  is tiny because then the likelihood surface becomes hyper-sensitive to  $\pi$ . The solution is to use the degenerate likelihood function for the likelihood calculation and the EM update equations. MARSS will implement this automatically when  $\mathbf{Q}$  or  $\mathbf{R}$  diagonal elements are set to zero and will try setting  $\mathbf{Q}$  and  $\mathbf{R}$  terms to zero automatically if `control$allow.degen=TRUE`. One odd case can occur when you set  $\mathbf{R}$  goes to zero (a matrix of zeros), but you are estimating  $\mathbf{x}_0$  (the prior for  $\mathbf{x}$ ). If `control$kf.x0="x10"` then  $\mathbf{x}_0$  is defined as the prior at  $t = 1$ . Then if  $\mathbf{x}_0 \equiv \mathbf{y}_1$  as  $\mathbf{R}$  goes to zero, and the log-likelihood goes to infinity. But if you set  $\mathbf{R} = 0$ , the log-likelihood will be finite. The reason is that  $\mathbf{R} \approx 0$  and  $\mathbf{R} = 0$  specify different likelihoods. In the first, the determinant  $|\mathbf{R}|$  will appear, and this goes to positive infinity as  $\mathbf{R}$  goes to zero. In the second case,  $\mathbf{R}$  does not appear in the likelihood and so the determinant of  $\mathbf{R}$  does not appear. If some elements of the diagonal of  $\mathbf{R}$  are going to zero, you should be suspect of the parameter estimates.

---

<sup>4</sup> An example of a  $\pi$  with shared values is  $\pi = \begin{bmatrix} a \\ a \end{bmatrix}$ .

Sometimes the structure of your data, e.g. one data value followed by a long string of missing values, is causing an odd spike in the likelihood at  $\mathbf{R} \approx 0$ . Try manually setting  $\mathbf{R}$  equal to zero to get the correct log-likelihood<sup>5</sup>

## 1.5 Other related packages

Packages that will do Kalman filtering and smoothing are very common, but packages that estimate MARSS models, especially constrained MARSS models, are much less common. The following are those with which we are familiar, however there are certainly more packages for estimating MARSS models in the engineering and economics of which we are unfamiliar. The MARSS package is unusual in that it uses an EM algorithm (sensu Shumway and Stoffer) for maximizing the likelihood as opposed to a Newton-esque method (esp. BFGS). The package is also unusual in that it allows you to specify the initial conditions at  $t = 0$  or  $t = 1$ , allows degenerate models (with diagonal elements of  $\mathbf{R}$  or  $\mathbf{Q}$  equal to zero, and allows one to specify a diverse array of models with shared and fixed parameter elements.

**DLM** DLM is an R package for fitting MARSS models. Our impression is that it is mainly Bayesian focused but it does allow MLE estimation via the `optim()` function. It now has a book also, *Dynamic Linear Models with R* by Petris et al.

**sspirs** `sspirs` an R package for fitting ARSS (univariate) models with Gaussian, Poisson and binomial error distributions.

**dse** `dse` (Dynamic Systems Estimation) is a R package for multivariate Gaussian state space models with a focus on ARMA models.

**SsfPack** `SsfPack` is a package for Ox/Splus that fits constrained multivariate Gaussian state space models using mainly (it seems) the BFGS algorithm but the newer versions support other types of maximization. `SsfPack` is very flexible and written in C to be fast. It has been used extensively on statistical finance problems and is optimized for dealing with large (financial) data sets. It is used and documented in *Time Series Analysis by State Space Methods* by Durbin and Koopman, *An Introduction to State Space Time Series Analysis* by Commandeur and Koopman, and *Statistical Algorithms for Models in State Space Form: SsfPack 3.0*, by Koopman, Shephard, and Doornik.

**Brodgar** The Brodgar software was developed by Alain Zuur to do (among many other things) dynamic factor analysis, which involves a special case of a MARSS model. The methods and many example analyses are given in *Analyzing Ecological Data* by Zuur, Ieno and Smith. This is the one

---

<sup>5</sup> The likelihood returned when  $\mathbf{R} \approx 0$  is not incorrect. It is just not the likelihood that you probably want. You want the likelihood where the  $\mathbf{R}$  term is dropped because it is zero.

package that we are aware of that also uses the same (or very similar) EM algorithm used in our MARSS package.

**eViews** eViews is a commercial economics software that will estimate at least some types of MARSS models.

**KFAS** The KFAS R package provides a fast Kalman filter and smoother. Examples in the package show how to fit MARSS models using the KFAS functions and the `optim()` function for maximization using Newton algorithms. The MARSS package uses the filter and smoother functions from the KFAS package when feasible; the KFAS package uses the initial condition set at  $t = 1$  and cannot be used when the initial condition is set at  $t = 0$ .





## Overview of the package functions

The MARSS package is object-based. It has two main types of objects: a model object (class `marssm`) and a maximum-likelihood fitted model object (class `marssMLE`). A `marssm` object specifies the structure of the model to be fitted. It is an *R* code version of the MARSS equation (Equation 1.1). A `marssMLE` object specifies both the model and the information necessary for fitting (initial conditions, controls, method). If the model has been fitted, the `marssMLE` object will also have the parameter estimates and (optionally) confidence intervals and bias.

### 2.1 The `MARSS()` function

The function `MARSS()` is an interface to the core fitting functions in the MARSS package. It allows a user to fit a MARSS model using a list to describe the model structure. It returns `marssm` and `marssMLE` objects which the user can later use in other functions, e.g. simulating or computing bootstrap confidence intervals.

`MLEobj=MARSS(data, constraint=list(), ..., fit=TRUE)` This function will fit a MARSS model to the data using a constraint list which is a list describing the structure of the model parameter matrices. The default model has a one-to-one correspondence between the state processes and observation time series (**Z** is the identity matrix). The default has a diagonal observation error matrix (**R**) and an unconstrained process-error matrix (**Q**). The output is a `marssMLE` object where the estimated parameter matrices are in `MLEobj$par`. If `fit=FALSE`, it returns a minimal `marssMLE` object that is ready for passing to a fitting function (below) but with no `par` element.

## 2.2 Core functions for fitting a MARSS model

The following core functions are designed to work with ‘unfitted’ `marssMLE` objects, that is a `marssMLE` object without the `par` element. Users do not normally need to call the `MARSSkem` or `MARSSoptim` functions since `MARSS()` will call those. Below `MLEobj` means the argument is a `marssMLE` object. Note, these functions can be called with a `marssMLE` object with a `par` element, but these functions will overwrite that element.

`MLEobj=MARSSkem(MLEobj)` This will fit a MARSS model via the Kalman-EM algorithm to the data using a properly specified `marssMLE` object, which has data, the `marssm` object and the necessary initial condition and control elements. See the appendix on the object structures in the MARSS package. `MARSSkem` does no error-checking. See `is.marssMLE().MARSSkem` uses `MARSSkf` described below.

`MLEobj=MARSSoptim(MLEobj)` This will fit a MARSS model via the BFGS algorithm provided in `optim()`. This requires a properly specified `marssMLE` object, such as would be passed to `MARSSkem`.

`MLEobj=MARSSmcinit(MLEobj)` This will perform a Monte Carlo initial conditions search and update the `marssMLE` object with the best initial conditions from the search.

`is.marssMLE(MLEobj)` This will check that a `marssMLE` object is properly specified and ready for fitting. This should be called before `MARSSkem` or `MARSSoptim` is called. This function is not typically needed if using `MARSS()` since `MARSS()` builds the model object for the user and does error-checking on model structure.

## 2.3 Functions for a fitted marssMLE object

The following functions use a `marssMLE` object that has a populated `par` element, i.e. a `marssMLE` object returned from one of the fitting functions (`MARSS`, `MARSSkem`, `MARSSoptim`). Below `modelObj` means the argument is a `marssm` object and `MLEobj` means the argument is a `marssMLE` object. Type `?function.name` to see information on function usage and examples.

`kf=MARSSkf(data, parList, ...)` This will compute the expected values of the hidden states given data via the Kalman filter (to produce estimates conditioned on  $1:t-1$ ) and the Kalman smoother (to produce estimates conditioned on  $1:T$ ). The function also returns the exact likelihood of the data conditioned on `parList`. A variety of other Kalman filter/smoothing information is also output (`kf` is a list of output); see `?MARSSkf` for more details.

`MLEobj=MARSSaic(MLEobj)` This adds model selection criteria, AIC, AICc, and AICb, to a `marssMLE` object.

**boot=**MARSSboot(MLEobj) This returns a list containing bootstrapped parameters and data via parametric or innovations bootstrapping.

**MLEobj=**MARSShessian(MLEobj) This adds a numerically estimated Hessian matrix to a marssMLE object.

**MLEobj=**MARSSparamCIs(MLEobj) This adds standard errors, confidence intervals, and bootstrap estimated bias for the maximum-likelihood parameters using bootstrapping or the Hessian to the passed in marssMLE object.

**sim.data=**MARSSsimulate(parList) This returns simulated data from a MARSS model specified via a list of parameter matrices in **parList** (this is a list with elements Q, R, U, etc). Typically one would pass in **parList** using the **par** element in an marssMLE object (i.e., **MLEobj\$par**), but you could also construct the list manually.

**paramVec=**MARSSvectorizeparam(MLEobj) This returns the estimated (and only the estimated) parameters as a vector. This is useful for storing the results of simulations and for writing functions that fit MARSS models using R's **optim** function.

**new.MLEobj=**MARSSvectorizeparam(MLEobj, paramVec) This will return a marssMLE object in which the estimated parameters (which are in **MLEobj\$par** along with the fixed values) are replaced with the values in **paramVec**.

## 2.4 Functions for marssm objects

**is.marssm(modelObj)** This will check that the free and fixed matrices in a marssm object are properly specified. This function is not typically needed if using **MARSS()** since **MARSS()** builds the marssm object for the user and does error-checking on model structure.

**summary(modelObj)** This will print the model parameter matrices showing the fixed values (in parentheses) and the location of the estimated elements. The estimated elements are shown as g1, g2, g3, ... which indicates which elements are shared (i.e., forced to have the same value). For example, an i.i.d. **R** matrix would appear as a diagonal matrix with just g1 on the diagonal.



---

## The MARSS() function

The `MARSS()` function is an interface to the core functions and allows users to fit MARSS models using a list to specify the model structure. The MARSS model takes the form

$$\mathbf{x}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{u} + \mathbf{w}_t, \text{ where } \mathbf{w}_t \sim \text{MVN}(0, \mathbf{Q}) \quad (3.1a)$$

$$\mathbf{y}_t = \mathbf{Z}\mathbf{x}_t + \mathbf{a} + \mathbf{v}_t, \text{ where } \mathbf{v}_t \sim \text{MVN}(0, \mathbf{R}) \quad (3.1b)$$

$$\mathbf{x}_1 \sim \text{MVN}(\boldsymbol{\pi}, \mathbf{V}_1) \quad (3.1c)$$

The  $\mathbf{y}$  is a  $n \times T$  matrix of observations and the  $\mathbf{x}$  is a  $m \times T$  matrix of hidden states. For example, a  $\mathbf{y}$  data matrix of 3 inputs ( $n = 3$ ) measured for 10 time steps would look like

$$\mathbf{y} = \begin{bmatrix} 1 & 2 & -99 & -99 & 3.2 & \dots & 8 \\ 2 & 5 & 3 & -99 & 5.1 & \dots & 5 \\ 1 & -99 & 2 & 2.2 & -99 & \dots & 7 \end{bmatrix}$$

where -99 denotes a missing value.  $\mathbf{x}$  might look like (here  $m = 2$ ):

$$\mathbf{x} = \begin{bmatrix} 0.8 & 2.2 & 3 & 2.8 & 3.2 & \dots & 7.1 \\ 1.5 & 2.5 & 2.5 & 2.2 & 3.1 & \dots & 6 \end{bmatrix}$$

In the process equation,  $\mathbf{B}$  is a  $m \times m$  matrix,  $\mathbf{u}$  and  $\mathbf{w}_t$  are  $m \times 1$  matrices, and  $\mathbf{Q}$  is a  $m \times m$  variance-covariance matrix for the  $\mathbf{w}_t$  process errors. In the observation equation,  $\mathbf{Z}$  is a  $n \times m$  design matrix of zeros and ones where the row sums equal one<sup>1</sup>.  $\mathbf{Z}$  specifies which observation time series,  $y_{i,1:T}$ , is associated with which hidden state process,  $x_{j,1:T}$ .  $\mathbf{a}$  and  $\mathbf{v}_t$  are  $n \times 1$  matrices, and  $\mathbf{R}$  is a  $n \times n$  variance-covariance matrix for the  $\mathbf{v}_t$  observation errors.

In the `MARSS()` function, the user specifies the model by passing in a parameter constraint list:

---

<sup>1</sup> In our examples,  $\mathbf{Z}$  is a design matrix but it can actually be any fixed matrix.

```
MARSS(data, constraint=list(Z=Z.constraint, B=B.constraint,
    U=U.constraint, Q=Q.constraint, A=A.constraint,
    R=R.constraint, x0=pi.constraint, V0=V1.constraint) )
```

`data` must be a  $n \times T$  matrix, that is time goes across columns. The argument `constraint` is a list, where the list elements for each model parameter can be a text string, vector of factors or matrix that specifies the form of the parameter. The defaults are

`Z.constraint="identity"` each  $y$  in  $\mathbf{y}$  corresponds to one  $x$  in  $\mathbf{x}$   
`B.constraint="identity"` no interactions between the  $x$ 's in  $\mathbf{x}$   
`U.constraint="unequal"` the  $u$ 's in  $\mathbf{u}$  are all different  
`Q.constraint="diagonal and unequal"` process errors are independent but have different variances  
`R.constraint="diagonal and equal"` the observations are i.i.d.  
`A.constraint="scaling"`  $\mathbf{a}$  is a set of scaling factors  
`pi.constraint="unequal"` all initial states are different  
`V0.constraint="zero"` the initial condition  $\mathbf{x}_1$  is fixed but unknown

Examples of the other possible `constraint` options for each parameter are listed below. We show the forms using  $m = 3$  (the number of hidden state processes) and  $n = 3$  (number of observation time series).

### 3.1 Process equation constraint options

#### 3.1.1 B constraints

In MARSS 1.0,  $\mathbf{B}$  must be fixed. There are two ways that  $\mathbf{B}$  can be fixed:

`B.constraint="identity"` The  $\mathbf{B}$  matrix is the identity matrix:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

`B.constraint=matrix(..., nrow=m, ncol=m)` Passing in a  $m \times m$  numeric matrix, means that  $\mathbf{B}$  is fixed to the values in the matrix. The matrix must be numeric and all eigenvalues must fall within the unit circle<sup>2</sup>. Using the string "zero", sets  $\mathbf{B} = 0$ .

#### 3.1.2 u constraints

`U.constraint="equal"` There is only one value in  $\mathbf{u}$ :

$$\begin{bmatrix} u \\ u \\ u \end{bmatrix}$$

---

<sup>2</sup> meaning `all(abs(eigen(B)$values)>=1)`.

`U.constraint="unequal"` or `U.constraint="unconstrained"` There are  $m$  values in  $\mathbf{u}$ :

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix}$$

`U.constraint=factor(c(...))` Here the constraint is specified as a length  $m$  character or numeric vector of class factor. The vector of factors specifies which values in  $\mathbf{u}$  are shared. `U.constraint=factor(c(1,1,2))` means that  $\mathbf{u}$  has the following structure:

$$\begin{bmatrix} u_1 \\ u_1 \\ u_2 \end{bmatrix}$$

There are two values in  $\mathbf{u}$  in this case. The factor levels can be either numeric or character. `c(1,1,2)` is the same as `c("north","north","south")`.

`U.constraint=matrix(..., nrow=m, ncol=1)` Passing in a  $m \times 1$  numeric matrix, means that  $\mathbf{u}$  is fixed to the values in the matrix. In MARSS 1.0,  $\mathbf{u}$  cannot vary in time, even if fixed.  $\mathbf{u}$  can be set to all zeros (a  $m \times 1$  matrix) by setting `U.constraint="zero"`; you might want to use this if you de-trended your data.

`U.constraint=matrix(..., nrow=m, ncol=1)` Passing in a  $m \times 1$  character matrix, means that all elements in  $\mathbf{u}$  are estimated. Those elements that share a name will be forced to be the same value.

You can also pass in a matrix with fixed values and NAs. For example, `matrix(c(0.1, NA, NA))`:

$$\begin{bmatrix} 0.1 \\ NA \\ NA \end{bmatrix}$$

The fixed values will be set to the fixed values and the elements specified by NAs will be estimated. The estimated elements are independent and not forced to be equal.

### 3.1.3 Q constraints

`Q.constraint="diagonal and equal"` There is only one process variance value in this case:

$$\begin{bmatrix} \sigma^2 & 0 & 0 \\ 0 & \sigma^2 & 0 \\ 0 & 0 & \sigma^2 \end{bmatrix}$$

`Q.constraint="diagonal and unequal"` There are  $m$  process variance values in this case:

$$\begin{bmatrix} \sigma_1^2 & 0 & 0 \\ 0 & \sigma_2^2 & 0 \\ 0 & 0 & \sigma_3^2 \end{bmatrix}$$

`Q.constraint="unconstrained"` There are values on the diagonal and the off-diagonals of **Q** and the variances and covariances are all different:

$$\begin{bmatrix} \sigma_1^2 & \sigma_{1,2} & \sigma_{1,3} \\ \sigma_{1,2} & \sigma_2^2 & \sigma_{2,3} \\ \sigma_{1,3} & \sigma_{2,3} & \sigma_3^2 \end{bmatrix}$$

There are  $m$  process variances and  $(m^2 - m)/2$  covariances in this case, so  $(m^2 + m)/2$  values to be estimated.

`Q.constraint="equalvarcov"` There is one process variance and one covariance:

$$\begin{bmatrix} \sigma^2 & \beta & \beta \\ \beta & \sigma^2 & \beta \\ \beta & \beta & \sigma^2 \end{bmatrix}$$

`Q.constraint=as.factor(c(...))` The constraint is specified as a length  $m$  character or numeric vector of class factor. This specifies that **Q** is diagonal and the vector of factors specifies which values on the diagonal are shared. For example, `Q.constraint=factor(c(2,1,2))` means that **Q** takes the form:

$$\begin{bmatrix} \sigma_2^2 & 0 & 0 \\ 0 & \sigma_1^2 & 0 \\ 0 & 0 & \sigma_2^2 \end{bmatrix}$$

`Q.constraint=factor(c(1,1,2))` means that **Q** takes the form:

$$\begin{bmatrix} \sigma_1^2 & 0 & 0 \\ 0 & \sigma_1^2 & 0 \\ 0 & 0 & \sigma_2^2 \end{bmatrix}$$

The factor levels can be either numeric or character. `c(1,1,2)` is the same as `c("north","north","south")`.

`Q.constraint=matrix(..., nrow=m, ncol=m)` Passing in a  $m \times m$  matrix, means that **Q** is fixed to the values in the matrix. The matrix must be numeric. Note if  $m = 1$ , you still need to wrap its value in `matrix()` so that its class is matrix.

You can also pass in a diagonal matrix with fixed values and NAs:

$$\begin{bmatrix} NA & 0 & 0 \\ 0 & 0.1 & 0 \\ 0 & 0 & NA \end{bmatrix}$$

The fixed values will be set to the fixed values and the NAs will be estimated; the estimates are independent and not forced to be equal. The matrix must be diagonal (0s on the off-diagonals).



### 3.1.4 $\pi$ constraints

This sets the constraints on the initial conditions,  $\mathbf{x}_1$ . `pi.constraint` has the following options:

`pi.constraint="equal"` There is only initial state value.

$$\begin{bmatrix} \pi \\ \pi \\ \pi \end{bmatrix}$$

Warning: specifying shared values in  $\pi$  will tend to produce ill-conditioned matrices in the algorithm and lead to numerical instability. Avoid using `pi.constraint="equal"` if possible.

`pi.constraint="unequal"` or `pi.constraint="unconstrained"` These are equivalent. There are  $m$  initial state parameters.

$$\begin{bmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \end{bmatrix}$$

`pi.constraint=factor(c(...))` The constraint is specified is a length  $m$  character or numeric vector of class `factor`. The vector of factors specifies which initial states have the same value. For example, `factor(c(1,1,2))` means that  $\pi$  takes the form:

$$\begin{bmatrix} \pi_1 \\ \pi_1 \\ \pi_2 \end{bmatrix}$$

There are two initial state parameters in this case. The factor levels can be either numeric or character. `c(1,1,2)` is the same as `c("n","n","s")`. Warning: specifying shared values in  $\pi$  will tend to produce ill-conditioned matrices in the algorithm and lead to numerical instability. Avoid if possible.

`pi.constraint=matrix(..., nrow=m, ncol=1)` Passing in a  $m \times 1$  numeric matrix, means that the initial states are fixed to the values in the matrix. You can set the initial states to zero by using `pi.constraint="zero"`.

You can also pass in a matrix with fixed values and NAs:

$$\begin{bmatrix} 10 \\ NA \\ NA \end{bmatrix}$$

The fixed values will be set to the fixed values and the values specified by NAs will be estimated.

### 3.1.5 $\mathbf{V}_1$ constraints

The initial state variance must be fixed. The default behavior is to treat  $\mathbf{x}_1$  as fixed but unknown<sup>3</sup>. In this case,  $\mathbf{V}_1=0$ . You can also set  $\mathbf{V}_1$  to a non-zero value but in that case  $\boldsymbol{\pi}$  must be fixed. This would translate to using a fixed prior for the initial states.

`V0.constraint=matrix(..., nrow=m, ncol=m)` Passing in a  $m \times m$  matrix, means that the initial state variance is fixed to the values in the matrix. The matrix must be numeric and be a proper variance-covariance matrix.

In general, unless a fixed prior is desired, users should just leave off `x0.constraint` and `V0.constraint` from the constraint list and use the default behavior which is  $\mathbf{x}_1$  treated as fixed but unknown.

## 3.2 Observation equation constraints

### 3.2.1 $\mathbf{Z}$ constraint

In the `MARSS()` function,  $\mathbf{Z}$  is normally a  $n \times m$  design matrix that specifies which  $x_i$  hidden state time series corresponds to which  $y_j$  time series<sup>4</sup>. In this case, each  $y_j$  time series (each row in  $\mathbf{y}$ ) corresponds to one and only one  $x_i$  time series (row in  $\mathbf{x}$ ). The  $\mathbf{Z}$  constraint is normally specified as a length  $n$  vector of class factor. The  $i$ -th element of this vector specifies which row in  $\mathbf{x}$  that the  $i$ -th row in  $\mathbf{y}$  corresponds to. Below are some examples of  $\mathbf{Z}$  matrices; see Chapter 6 and the case studies chapters for more examples.

`Z.constraint=factor(c(1,1,1))` All  $\mathbf{y}$  time series are observing the same (and only) hidden state trajectory  $x$  ( $n = 3$  and  $m = 1$ ):

$$\mathbf{Z} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

`Z.constraint=factor(c(1,2,3))` Each time series in  $\mathbf{y}$  corresponds to a different hidden state trajectory. This is the default  $\mathbf{Z}$  constraint and in this case  $n = m$ :

$$\mathbf{Z} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

<sup>3</sup> Shumway and Stoffer use  $\mathbf{x}$  at  $t = 0$  as the initial state, but we follow Ghahramani's approach and use  $\mathbf{x}$  at  $t = 1$  as the initial state. Both approaches give the same answer but the EM update equations are slightly different.

<sup>4</sup> However, you can pass in other fixed but non-design  $\mathbf{Z}$  matrices.

`Z.constraint=factor(c(1,1,2))` The first two time series in  $\mathbf{y}$  corresponds to one hidden state trajectory and the third  $\mathbf{y}$  time series corresponds to a different hidden state trajectory. Here  $n = 3$  and  $m = 2$ :

$$\mathbf{Z} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}$$

The  $\mathbf{Z}$  constraint can be specified using either numeric or character factor levels. `c(1,1,2)` is the same as `c("north","north","south")`

`Z.constraint="identity"` This is the default behavior. This means  $\mathbf{Z}$  is a  $n \times n$  identity matrix and  $m = n$ . If  $n = 3$ , it is the same as `Z.constraint=factor(c(1,2,3))`.

`Z.constraint=matrix(..., nrow=n, ncol=m)` Passing in a  $n \times m$  numeric matrix, means that  $\mathbf{Z}$  is fixed to the values in the matrix. The matrix must be numeric but it does not need to be a design matrix.

### 3.2.2 a constraint

Using the `MARSS()` function, only `A.constraint="scaling"` or a fixed  $\mathbf{a}$  are allowed. If “scaling”,  $\mathbf{a}$  is a scaling factor without any direct meaning (it just says how the time series in  $\mathbf{y}$  scale relative to each other). Unless you know the correct scaling, because say you simulated the data or you de-meanded the data, you should use the default. To set  $\mathbf{a}$  to a fixed value (not estimated), use `A.constraint=matrix(..., nrow=n, ncol=m)`. The matrix must be numeric. To fix  $\mathbf{a}$  to zero, use `A.constraint="zero"`.

Note, you can circumvent the restriction on  $\mathbf{a}$  by passing in a matrix with NAs (for estimated elements) or passing in a `fixed/free` pair. But be aware that it is very easy to make your model unsolvable (an infinite number of solutions)—if for example you try to estimate more than  $n_i - 1$   $a$ ’s, where  $n_i$  is the number of  $y$ ’s for one  $x$ . If  $\mathbf{Z}$  is identity, then  $n_i = 1$  and you can estimate no  $a$ ’s; all must be fixed.

### 3.2.3 R constraint

The  $\mathbf{R}$  constraint is completely analogous to the  $\mathbf{Q}$  constraint, except that it is  $n \times n$  instead of  $m \times m$ . Its allowable constraints are affected by the presence of missing data points in  $\mathbf{y}$ . If data are missing, then  $\mathbf{R}$  must be diagonal.



## Model specification in the core functions

Most users will not directly work with the core functions nor build `marssm` objects from scratch. Instead, they will usually interact with the core functions via the function `MARSS()` described in Chapter 3. With the `MARSS()` function, the user specifies the model structure with text strings (“diagonal”, “unconstrained”, etc.) and `MARSS()` builds the `marssm` object. However, a basic understanding of the structure of `marssm` objects is necessary if one wants to fit more flexible models or to interact directly with the core functions.

The first step of model specification is to write down (e.g. on paper) the model in matrix form (Equation 1.1) with notes on the dimensions (rows and columns) of each parameter and of  $\mathbf{x}$  and  $\mathbf{y}$ . In the core functions, the parameters in the `MARSS` model must be passed as matrices of the correct dimension, and the parameters in the R functions correspond one-to-one to the mathematical equation. For example,  $\mathbf{u}$  must be passed in as a matrix of dimension `c(m,1)`. The function will return an error if anything else is passed in (including a matrix with `dim=c(1,m)` or a vector of length  $m$ ).

### 4.1 The fixed and free components of the model parameters

In a `marssm` object, each parameter must be specified by a pair of matrices: `free` which gives the location and sharing of the estimated elements in the parameter matrix and `fixed` which specifies the location and value of the fixed elements in the parameter matrix. For example,  $\mathbf{Q}$  is specified by `free$Q` and `fixed$Q`.

#### 4.1.1 The fixed matrices

The fixed matrix specifies the values (numeric) of the fixed (meaning not estimated) elements. In the fixed matrix, the free (meaning estimated or fitted)

elements are denoted with `NA`. The following shows some common examples of the fixed matrix using `fixed$Q` as the example. Each of the other fixed matrices for the other parameters uses the same pattern.

`Q` is unconstrained, so there are no fixed values

$$\text{fixed\$Q} = \begin{bmatrix} NA & NA & NA \\ NA & NA & NA \\ NA & NA & NA \end{bmatrix}$$

`Q` is a diagonal matrix, so the off-diagonals are fixed at 0. The diagonal elements will be estimated.

$$\text{fixed\$Q} = \begin{bmatrix} NA & 0 & 0 \\ 0 & NA & 0 \\ 0 & 0 & NA \end{bmatrix}$$

`Q` is fixed, i.e. will not be estimated rather all values in the `Q` matrix are fixed.

$$\text{fixed\$Q} = \begin{bmatrix} 0.1 & 0 & 0 \\ 0 & 0.1 & 0 \\ 0 & 0 & 0.1 \end{bmatrix}$$

#### 4.1.2 The free matrices

The free matrix specifies which elements are estimated and specifies how (and whether) the free elements are shared. In the free matrix, the fixed elements are denoted `NA`. The following shows some common examples of `free` using `free$Q` as the example. `free` can be passed in as a character matrix or a numeric matrix, but if numeric, the numeric will be changed to character (thus 0 becomes “0” and is the name “0” not the number 0).

`Q` is a diagonal matrix in which there is only one, shared, value on the diagonal. Thus there is only one `Q` parameter.

$$\text{free\$Q} = \begin{bmatrix} 1 & NA & NA \\ NA & 1 & NA \\ NA & NA & 1 \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} \text{“a”} & NA & NA \\ NA & \text{“a”} & NA \\ NA & NA & \text{“a”} \end{bmatrix}$$

Here “1” does not mean “number 1” but rather that the name of the shared parameter is “1”. On the example at the right, the name of the shared parameter is “a”.

`Q` is a diagonal matrix in which each of the diagonal elements are different.

$$\text{free\$Q} = \begin{bmatrix} 1 & NA & NA \\ NA & 2 & NA \\ NA & NA & 3 \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} \text{“north”} & NA & NA \\ NA & \text{“middle”} & NA \\ NA & NA & \text{“south”} \end{bmatrix}$$

$\mathbf{Q}$  has one value on the diagonal and another one on the off-diagonals. There are no fixed values in  $\mathbf{Q}$ .

$$\text{free}\$Q = \begin{bmatrix} 1 & 2 & 2 \\ 2 & 1 & 2 \\ 2 & 2 & 1 \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} "a" & "b" & "b" \\ "b" & "a" & "b" \\ "b" & "b" & "a" \end{bmatrix}$$

$\mathbf{Q}$  is unconstrained. There are no fixed values in  $\mathbf{Q}$  in this case. Note that since,  $\mathbf{Q}$  is a variance-covariance matrix, it must be symmetric across the diagonal.

$$\text{free}\$Q = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 5 \\ 3 & 5 & 6 \end{bmatrix}$$

## 4.2 Limits on the model forms that can be fit

MARSS version 1.0 will allow any combination of fixed and shared values for  $\mathbf{a}$  and  $\mathbf{u}$ , but for  $\mathbf{R}$ ,  $\mathbf{Q}$ ,  $\mathbf{B}$ , and  $\mathbf{Z}$  there are limits to what forms these matrices can take. These limitations have to do with the way the EM algorithm is coded for version 1.0. Version 2.0 will remove many of these restrictions.

- $\mathbf{R}$  and  $\mathbf{Q}$  can be fixed, unconstrained, diagonal with any pattern of fixed and shared values on the diagonal, or a matrix with one estimated value on the diagonal and another estimated value for the off-diagonals.
- If there are missing values in the data,  $\mathbf{R}$  must be diagonal (or fixed).
- $\mathbf{B}$  cannot be estimated. It must be fixed but it need not be diagonal. However, all its eigen values must fall on the unit circle<sup>1</sup>.
- $\mathbf{Z}$  cannot be estimated. It must be fixed. Although in all our examples,  $\mathbf{Z}$  is a design matrix (meaning only 0s and 1s and all row sums equal to 1),  $\mathbf{Z}$  is not required to be a design matrix.

The other limitation is that one must specify a model that has only one solution. The core MARSS functions will allow you to attempt to fit an improper model (one with multiple solutions). If you do this accidentally, it may or may not be obvious that you have a problem. The MARSS estimation functions may chug along happily and return a solution. Careful thought about your model structure and the structure of the estimated parameter matrices will help you determine if your model is under-constrained and unsolvable. Basically, take care when using MARSS core functions directly and remember that they will not prevent you from fitting an under-constrained model. This is not a problem when using the function `MARSS()`. The `MARSS()` function includes error-checks that are designed to prevent users from specifying an under-constrained models (for example it only allows  $\mathbf{a}$  to be fixed or act like a scaling factor).

<sup>1</sup> This means that the absolute value of all the eigenvalues must be less than or equal to 1 (`all(abs(eigen(B)$values)<=1)`).





---

## Algorithms used in the MARSS package

### 5.1 Kalman filter and smoother

The MARSS model (Equation 1.1) is a linear dynamical system in discrete time. In 1960, Rudolf Kalman published the Kalman filter (Kalman, 1960), a recursive algorithm that solves for the expected value of the hidden state(s) at time  $t$  conditioned on the data up to time  $t$ :  $E(\mathbf{X}_t | \mathbf{y}_1^t)$ . The Kalman filter gives the optimal (lowest mean square error) estimate of the unobserved  $\mathbf{x}_t$  based on the observed data up to time  $t$  for this class of linear dynamical system. The Kalman smoother (Rauch et al., 1965) solves for the expected value of the hidden state(s) conditioned on all the data:  $E(\mathbf{X}_t | \mathbf{y}_1^T)$ . If the errors in the stochastic process are Gaussian, then the estimators from the Kalman filter and smoother are also the maximum-likelihood estimates.

However, even if the errors are not Gaussian, the estimators are optimal in the sense that they are estimators with the least variability possible. This robustness is one reason the Kalman filter is so powerful—it provides well-behaving estimates of the hidden states for all kinds of multivariate autoregressive processes, not just Gaussian processes. The Kalman filter and smoother are widely used in time-series analysis, and there are many textbooks covering it and its applications. In the interest of giving the reader a single point of reference, we use Shumway and Stoffer (2006) as our reference and adopt their notation (for the most part). The algorithms used in MARSS are not computationally efficient; see Koopman et al. (1999, sec. 4.3) for a more efficient Kalman filter implementation.

The `MARSSkf` function provides the Kalman filter and smoother and has the following outputs:

- `x tt1` The expected value of  $\mathbf{X}_t$  conditioned on the data up to time  $t - 1$ .
- `x tt` The expected value of  $\mathbf{X}_t$  conditioned on the data up to time  $t$ .
- `x tT` The expected value of  $\mathbf{X}_t$  conditioned on all the data from time 1 to  $T$ .  
This is the smoothed state estimate.

- Vtt1** The variance of  $\mathbf{X}_t$  conditioned on the data up to time  $t - 1$ . Denoted  $P_t^{t-1}$  in section 6.2 in Shumway and Stoffer (2006).
- Vtt** The variance of  $\mathbf{X}_t$  conditioned on the data up to time  $t$ . Denoted  $P_t^t$  in section 6.2 in Shumway and Stoffer (2006).
- VtT** The variance of  $\mathbf{X}_t$  conditioned on all the data from time 1 to  $T$ .
- Vtt1T** The covariance of  $\mathbf{X}_t$  and  $\mathbf{X}_{t-1}$  conditioned on all the data, 1 to  $T$ .
- Kt** The Kalman gain. This is part of the update equations and relates to the amount **xtt1** is updated by the data at time  $t$  to produce **xtt**.
- J** This is similar to the Kalman gain but is part of the Kalman smoother. See Equation 6.49 in Shumway and Stoffer (2006).
- Innov** This has the innovations at time  $t$ , defined as  $\epsilon_t \equiv \mathbf{y}_t - \mathbf{E}(\mathbf{Y}_t)$ . These are the residuals, the difference between the data and their predicted values. See Equation 6.24 in Shumway and Stoffer (2006).
- Sigma** This has the  $\Sigma_t$ , the variance-covariance matrices for the innovations at time  $t$ . This is used for the calculation of confidence intervals, the s.e. on the state estimates and the likelihood. See Equation 6.25 in Shumway and Stoffer (2006) for the  $\Sigma_t$  calculation.
- logLik** The log-likelihood of the data conditioned on the model parameters. See the section below on the likelihood calculation.

## 5.2 The exact likelihood

The likelihood of the data given a set of MARSS parameters is part of the output of the **MARSSkf** function. The likelihood computation is based on the innovations form of the likelihood (Schweppe, 1965) and uses the output from the Kalman filter:

$$\log L(\Theta|data) = -\frac{N}{2\log(2\pi)} - \frac{1}{2} \left( \sum_{t=1}^T \log |\Sigma_t| + \sum_{t=1}^T (\epsilon_t)^\top \Sigma_t^{-1} \epsilon_t \right) \quad (5.1)$$

where  $N$  is the total number of data points,  $\epsilon_t$  is the innovations at time  $t$  and  $|\Sigma_t|$  is the determinant of the innovations variance-covariance matrix at time  $t$ . Reference Equation 6.62 in Shumway and Stoffer (2006). However there are a few differences between the log-likelihood output by **MARSSkf** and that described in Shumway and Stoffer (2006).

The standard likelihood calculation (Equation 6.62 in Shumway and Stoffer (2006)) is biased when there are missing values in the data, and the missing data modifications discussed in Section 6.4 in Shumway and Stoffer (2006) do not correct for this bias. Harvey (1989), Section 3.4.7, discusses at length that the standard missing values correction leads to an inexact likelihood when there are missing values. The bias is minor if there are few missing values, but it becomes severe as the number of missing values increases. Many ecological datasets may have over 25% missing values and this level of missing values

leads to a very biased likelihood if one uses the inexact formula. Harvey (1989) provides some non-trivial ways to compute the exact likelihood.

We use instead the exact likelihood correction for missing values that is presented in Section 12.3 in Brockwell and Davis (1991). This solution is straight-forward to implement. The correction involves the following changes to  $\epsilon_t$  and  $\Sigma_t$  in the Equation 5.1. Suppose the value  $y_{i,t}$  is missing. First, the corresponding  $i$ -th value of  $\epsilon_t$  is set to 0. Second, the  $i$ -th diagonal value of  $\Sigma_t$  is set to 1 and the off-diagonal elements on the  $i$ -th column and  $i$ -th row are set to 0.

## 5.3 Maximum-likelihood parameter estimation

### 5.3.1 Kalman-EM algorithm

The MARSS package provides a maximum-likelihood algorithm which uses an Expectation-Maximization (EM) algorithm (function `MARSSkem`) with the Kalman smoother. EM algorithms are widely used algorithms that extend maximum-likelihood estimation to cases where there are hidden random variables in a model (Dempster et al., 1977; Harvey, 1989; Harvey and Shephard, 1993; McLachlan and Krishnan, 2008).

The EM algorithm finds the maximum-likelihood estimates of the parameters in a MARSS model using an iterative process. Starting with an initial set of parameters<sup>1</sup>, which we will denote  $\hat{\Theta}_1$ , an updated parameter set  $\hat{\Theta}_2$  is obtained by finding the  $\hat{\Theta}_2$  that maximizes the expected value of the likelihood over the distribution of the states ( $\mathbf{X}$ ) conditioned on  $\hat{\Theta}_1$ :

$$\hat{\Theta}_2 = \arg \max_{\Theta} E_{\mathbf{X}|\hat{\Theta}_1} [\log L(\Theta | \mathbf{Y}_1^T = \mathbf{y}_1^T, \mathbf{X})] \quad (5.2)$$

Then using  $\hat{\Theta}_2$  in place of  $\hat{\Theta}_1$  in Equation (5.2), an updated parameter set  $\hat{\Theta}_3$  is calculated. This is repeated until the expected log-likelihood stops increasing (or increases less than some set tolerance level).

Implementing this algorithm is straight-forward, hence its popularity.

1. Set an initial set of parameters,  $\hat{\Theta}_1$
2. E step: using the model for the hidden states ( $\mathbf{X}$ ) and  $\hat{\Theta}_1$ , calculate the expected values of  $\mathbf{X}$  conditioned on all the data  $\mathbf{y}_1^T$ ; this is `xtT` output by `MARSSkf`. Also calculate expected values of any functions of  $\mathbf{X}$ ,  $g(\mathbf{X})$ , that appear in your expected log-likelihood function.
3. M step: put those  $E(\mathbf{X} | \mathbf{Y}_1^T = \mathbf{y}_1^T, \hat{\Theta}_1)$  and  $E(g(\mathbf{X}) | \mathbf{Y}_1^T = \mathbf{y}_1^T, \hat{\Theta}_1)$  into your expected log-likelihood function in place of  $\mathbf{X}$  (and  $g(\mathbf{X})$ ) and maximize with respect to  $\Theta$ . This gives you  $\hat{\Theta}_2$ .

---

<sup>1</sup> You can choose these however you wish, however choosing something not too far off from the correct values will make the algorithm go faster.

4. Repeat the E and M steps until the log likelihood stops increasing.

The EM equations in our algorithm, which we term the Kalman-EM algorithm, are extensions of those in Shumway and Stoffer (1982) and Ghahramani and Hinton (1996). Our Kalman-EM algorithm is an extended version because our algorithm is for cases where there are constraints within the parameter matrices (shared values, diagonal structure, block-diagonal structure, ...) and where there are fixed values within the parameter matrices. Holmes (2010) gives the full derivation of our EM algorithm.

The EM algorithm is a hill-climbing algorithm and like all hill-climbing algorithms can get stuck on local maxima. The MARSS package includes a Monte-Carlo initial conditions searcher (function `MARSSmcinit`) based on Biernacki et al. (2003) to minimize this problem. EM algorithms are also known to get close to the maximum very quickly but then creep toward the absolute maximum. Quasi-Newton methods find the absolute maximum much faster, but they can be sensitive to initial conditions. We have found that with MARSS models, quasi-Newton methods (at least using `optim`) will sometimes converge far from the maximum even when started close to the known maximum. For this reason, the Monte Carlo initial condition search that works for EM algorithms may not work for Newton algorithms.

## 5.4 Parametric and innovations bootstrapping

Bootstrapping can be used to construct frequentist confidence intervals on the parameter estimates (Stoffer and Wall, 1991) and to compute the small-sample AIC corrector for MARSS models (Cavanaugh and Shumway, 1997); the functions `MARSSparamCIs` and `MARSSaic` do these computations.

The `MARSSboot` function provides both parametric and innovations bootstrapping of MARSS models. The innovations bootstrap algorithm by Stoffer and Wall (1991) bootstraps the model residuals (the innovations). This is a semi-parametric bootstrap since it uses, partially, the maximum-likelihood parameter estimates. This algorithm cannot be used if there are missing values in the data. Also for short time series, it gives biased bootstraps because one cannot resample the first few innovations.

`MARSSboot` also provides a fully parametric bootstrap. This uses the maximum-likelihood MARSS parameters to simulate data from which bootstrap parameter estimates are obtained. Our research (Holmes and Ward, 2010) indicates that this provides unbiased bootstrap parameter estimates, and it works with datasets with missing values. Lastly, `MARSSboot` also outputs parameters sampled from a numerically estimated Hessian matrix.

## 5.5 Simulation and forecasting

The `MARSSsimulate` function simulates from a MARSS model using a list of parameter matrices. It use the `mvrnorm` function to produce draws of the process and observation errors from multivariate normal distributions for each time step.

## 5.6 Model selection

The package provides a `MARSSaic` function for computing AIC, AICc and AICb. The latter is a small-sample corrector for autoregressive state-space models. The bias problem with AIC and AICc for short time-series data has been shown in Cavanaugh and Shumway (1997) and Holmes and Ward (2010). AIC and AICc tend to select overly complex MARSS models when the time-series data are short. AICb corrects this bias. The algorithm for a non-parametric AICb is given in Cavanaugh and Shumway (1997). Their algorithm uses the innovations bootstrap (Stoffer and Wall, 1991), which means it cannot be used when there are missing data. We added a parametric AICb (Holmes and Ward, 2010), which uses a parametric bootstrap. This algorithm allows one to compute AICb when there are missing data and it provides unbiased AIC even for short time series. See Holmes and Ward (2010) for discussion and testing of parametric AICb for MARSS models.

AICb is comprised of the familiar AIC fit term,  $-2 \log L$ , plus a penalty term that is the mean difference between the log likelihood the data under the bootstrapped maximum-likelihood parameter estimates and the log likelihood of the data under the original maximum-likelihood parameter estimate:

$$AICb = -2 \log L(\hat{\theta}|\mathbf{y}) + 2 \left( \frac{1}{N_b} \sum_{i=1}^{N_b} -\log \frac{L(\hat{\theta}^*(i)|\mathbf{y})}{L(\hat{\theta}|\mathbf{y})} \right) \quad (5.3)$$

where  $\hat{\theta}$  is the maximum-likelihood parameter set under the original data  $\mathbf{y}$ ,  $\hat{\theta}^*(i)$  is a maximum-likelihood parameter set estimated from the  $i$ -th bootstrapped data set  $\mathbf{y}^*(i)$ , and  $N_b$  is the number of bootstrap data sets. It is important to notice that the likelihood in the AICb equation is  $L(\hat{\theta}^*|\mathbf{y})$  not  $L(\hat{\theta}^*|\mathbf{y}^*)$ . In other words, we are taking the average of the likelihood of the original data given the bootstrapped parameter sets.



## Examples

Here we show a series of short examples using the MARSS package functions. Chapters 8–12 give case studies which walk through detailed multi-level analyses. The examples in this chapter use the Washington harbor seal dataset (`?harborSealWA`), which has five observation time series. First set up the data to make time go across the columns and to remove the year column:

```
dat = t(harborSealWA)
dat = dat[2:nrow(dat),] #remove the year row
```

### 6.1 Fitting different MARSS models to a dataset

#### 6.1.1 One hidden state process for each observation time series

This is the default model for the `MARSS()` function. In this case,  $n = m$ , the observation errors are i.i.d. and the process errors are independent and have different variances. The elements in  $\mathbf{u}$  are all different (meaning, they are not forced to be the same). Mathematically, the MARSS model being fit is:

$$\begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \\ x_{4,t} \\ x_{5,t} \end{bmatrix} = \begin{bmatrix} x_{1,t-1} \\ x_{2,t-1} \\ x_{3,t-1} \\ x_{4,t-1} \\ x_{5,t-1} \end{bmatrix} + \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ u_5 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix}, \mathbf{w}_t \sim \text{MVN} \left( 0, \begin{bmatrix} q_1 & 0 & 0 & 0 & 0 \\ 0 & q_2 & 0 & 0 & 0 \\ 0 & 0 & q_3 & 0 & 0 \\ 0 & 0 & 0 & q_4 & 0 \\ 0 & 0 & 0 & 0 & q_5 \end{bmatrix} \right)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \\ x_{4,t} \\ x_{5,t} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} v_{1,t} \\ v_{2,t} \\ v_{3,t} \\ v_{4,t} \\ v_{5,t} \end{bmatrix}, \mathbf{v}_t \sim \text{MVN} \left( 0, \begin{bmatrix} r & 0 & 0 & 0 & 0 \\ 0 & r & 0 & 0 & 0 \\ 0 & 0 & r & 0 & 0 \\ 0 & 0 & 0 & r & 0 \\ 0 & 0 & 0 & 0 & r \end{bmatrix} \right)$$

To fit this model, use `MARSS()`. The function will output the basic model structure and the time to convergence. Notice that the output warns you that the convergence tolerance is high. You can set it lower by passing in `control=list(conv.test.slope.tol=0.1)`.

```
kemfit = MARSS(dat)
```

Success! Parameters converged at 30 iterations.

Alert: `conv.test.slope.tol` is 0.5.

Test with smaller values (<0.1) to ensure convergence.

MARSS fit is

Estimation method: kem

Estimation converged in 30 iterations.

Log-likelihood: 19.11268

AIC: -6.225352 AICc: 3.848722

	Estimate
Q.1	0.03024
Q.2	0.01072
Q.3	0.00690
Q.4	0.00426
Q.5	0.05267
R.1	0.00939
U.1	0.06840
U.2	0.07173
U.3	0.04169
U.4	0.05228
U.5	-0.00288
x0.1	5.98907
x0.2	6.72825
x0.3	6.66751
x0.4	5.83982
x0.5	6.60549

Standard errors have not been calculated.

Use `MARSSparamCIs` to compute CIs and bias estimates.

`MARSS()` is automatically creating parameter names (since you did not tell it the names). To see exactly where each parameter element (e.g. Q.2) appears in its parameter matrix, type `summary(kemfit$model)`.

The default method is the Kalman-EM algorithm (`method="kem"`). You can use a quasi-Newton method (BFGS) by setting `method="BFGS"`.

```
kemfit.bfgs = MARSS(dat, method="BFGS")
```

Success! Converged in 100 iterations.



```

MARSS fit is
Estimation method: BFGS
Estimation converged in 100 iterations.
Log-likelihood: 19.13936
AIC: -6.278712   AICc: 3.795362

```

	Estimate
Q.1	0.03368
Q.2	0.01124
Q.3	0.00722
Q.4	0.00437
Q.5	0.05600
R.1	0.00849
U.1	0.06837
U.2	0.07152
U.3	0.04187
U.4	0.05233
U.5	-0.00270
x0.1	5.98432
x0.2	6.72169
x0.3	6.65689
x0.4	5.83529
x0.5	6.60423

Standard errors have not been calculated.  
 Use `MARSSparamCIs` to compute CIs and bias estimates.

Using the default EM convergence criteria, the EM algorithm stops at a log-likelihood a little lower than the BFGS algorithm does, but the EM algorithm was considerable faster, 24 times faster in this case. If you wanted to use the Kalman-EM fit as the initial conditions, pass in the `inits` argument.

```
kemfit.bfgs2 = MARSS(dat, method="BFGS", inits=kemfit$par)
```

Output not shown.

### 6.1.2 Five correlated hidden state processes

This is the same model except that the hidden states have temporally correlated process errors. Mathematically, this is the model:

$$\begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \\ x_{4,t} \\ x_{5,t} \end{bmatrix} = \begin{bmatrix} x_{1,t-1} \\ x_{2,t-1} \\ x_{3,t-1} \\ x_{4,t-1} \\ x_{5,t-1} \end{bmatrix} + \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ u_5 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix}, \mathbf{w}_t \sim \text{MVN} \left( 0, \begin{bmatrix} q_1 & c_{1,2} & c_{1,3} & c_{1,4} & c_{1,5} \\ c_{1,2} & q_2 & c_{2,3} & c_{2,4} & c_{2,5} \\ c_{1,3} & c_{2,3} & q_3 & c_{3,4} & c_{3,5} \\ c_{1,4} & c_{2,4} & c_{3,4} & q_4 & c_{4,5} \\ c_{1,5} & c_{2,5} & c_{3,5} & c_{4,5} & q_5 \end{bmatrix} \right)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \\ x_{4,t} \\ x_{5,t} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} v_{1,t} \\ v_{2,t} \\ v_{3,t} \\ v_{4,t} \\ v_{5,t} \end{bmatrix}, \mathbf{v}_t \sim \text{MVN} \left( 0, \begin{bmatrix} r & 0 & 0 & 0 & 0 \\ 0 & r & 0 & 0 & 0 \\ 0 & 0 & r & 0 & 0 \\ 0 & 0 & 0 & r & 0 \\ 0 & 0 & 0 & 0 & r \end{bmatrix} \right)$$

To fit, use `MARSS()` with the constraint argument set. The output is not shown but it will appear if you type this on the R command line.

```
kemfit = MARSS(dat, constraint=list(Q="unconstrained"))
```

### 6.1.3 Five equally correlated hidden state processes

This is the same model except that now there is only one process error variance and one process error covariance. Mathematically, the model is:

$$\begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \\ x_{4,t} \\ x_{5,t} \end{bmatrix} = \begin{bmatrix} x_{1,t-1} \\ x_{2,t-1} \\ x_{3,t-1} \\ x_{4,t-1} \\ x_{5,t-1} \end{bmatrix} + \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ u_5 \end{bmatrix} + \begin{bmatrix} v_{1,t} \\ v_{2,t} \\ v_{3,t} \\ v_{4,t} \\ v_{5,t} \end{bmatrix}, \mathbf{e}_t \sim \text{MVN} \left( 0, \begin{bmatrix} q & c & c & c & c \\ c & q & c & c & c \\ c & c & q & c & c \\ c & c & c & q & c \\ c & c & c & c & q \end{bmatrix} \right)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \\ x_{4,t} \\ x_{5,t} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix}, \mathbf{w}_t \sim \text{MVN} \left( 0, \begin{bmatrix} r & 0 & 0 & 0 & 0 \\ 0 & r & 0 & 0 & 0 \\ 0 & 0 & r & 0 & 0 \\ 0 & 0 & 0 & r & 0 \\ 0 & 0 & 0 & 0 & r \end{bmatrix} \right)$$

To fit, use the following code (output not shown):

```
kemfit = MARSS(dat, constraint=list(Q="equalvarcov"))
```

### 6.1.4 Five hidden state processes with a “north” and a “south” $u$ parameter

Here we fit a model with five independent hidden states where each observation time series is an independent observation of a different hidden trajectory

but the hidden trajectories 1-3 share their  $u$  and  $q$  parameters, while hidden trajectories 4-5 share theirs. This is the model:

$$\begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \\ x_{4,t} \\ x_{5,t} \end{bmatrix} = \begin{bmatrix} x_{1,t-1} \\ x_{2,t-1} \\ x_{3,t-1} \\ x_{4,t-1} \\ x_{5,t-1} \end{bmatrix} + \begin{bmatrix} u_n \\ u_n \\ u_n \\ u_s \\ u_s \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix}, \mathbf{w}_t \sim \text{MVN} \left( 0, \begin{bmatrix} q_n & 0 & 0 & 0 & 0 \\ 0 & q_n & 0 & 0 & 0 \\ 0 & 0 & q_n & 0 & 0 \\ 0 & 0 & 0 & q_s & 0 \\ 0 & 0 & 0 & 0 & q_s \end{bmatrix} \right)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \\ x_{4,t} \\ x_{5,t} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} v_{1,t} \\ v_{2,t} \\ v_{3,t} \\ v_{4,t} \\ v_{5,t} \end{bmatrix}, \mathbf{v}_t \sim \text{MVN} \left( 0, \begin{bmatrix} r & 0 & 0 & 0 & 0 \\ 0 & r & 0 & 0 & 0 \\ 0 & 0 & r & 0 & 0 \\ 0 & 0 & 0 & r & 0 \\ 0 & 0 & 0 & 0 & r \end{bmatrix} \right)$$

To fit use the following code (output not shown):

```
regions=factor(c("N","N","N","S","S"))
kemfit = MARSS(dat, constraint=list(U=regions, Q=regions))
```

### 6.1.5 Fixed observation error variance

Here we fit the same model but with a known observation error variance. This is the model:

$$\begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \\ x_{4,t} \\ x_{5,t} \end{bmatrix} = \begin{bmatrix} x_{1,t-1} \\ x_{2,t-1} \\ x_{3,t-1} \\ x_{4,t-1} \\ x_{5,t-1} \end{bmatrix} + \begin{bmatrix} u_n \\ u_n \\ u_n \\ u_s \\ u_s \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix}, \mathbf{w}_t \sim \text{MVN} \left( 0, \begin{bmatrix} q_n & 0 & 0 & 0 & 0 \\ 0 & q_n & 0 & 0 & 0 \\ 0 & 0 & q_n & 0 & 0 \\ 0 & 0 & 0 & q_s & 0 \\ 0 & 0 & 0 & 0 & q_s \end{bmatrix} \right)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \\ x_{4,t} \\ x_{5,t} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} v_{1,t} \\ v_{2,t} \\ v_{3,t} \\ v_{4,t} \\ v_{5,t} \end{bmatrix},$$

$$\mathbf{v}_t \sim \text{MVN} \left( 0, \begin{bmatrix} 0.01 & 0 & 0 & 0 & 0 \\ 0 & 0.01 & 0 & 0 & 0 \\ 0 & 0 & 0.01 & 0 & 0 \\ 0 & 0 & 0 & 0.01 & 0 \\ 0 & 0 & 0 & 0 & 0.01 \end{bmatrix} \right)$$

To fit this model, use the following code (output not shown):

```
regions=factor(c("N","N","N","S","S"))
kemfit = MARSS(dat, constraint=list(U=regions, Q=regions,
                                     R=diag(0.01,5)))
```

### 6.1.6 One hidden state and five i.i.d. observation time series

Instead of five hidden state trajectories, we specify that there is only one and all the observations are of that one trajectory. Mathematically, the model is:

$$x_t = x_{t-1} + u + v_t, \quad v_t \sim N(0, q)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} x_t + \begin{bmatrix} 0 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix}, \quad \mathbf{w}_t \sim \text{MVN} \left( 0, \begin{bmatrix} r & 0 & 0 & 0 & 0 \\ 0 & r & 0 & 0 & 0 \\ 0 & 0 & r & 0 & 0 \\ 0 & 0 & 0 & r & 0 \\ 0 & 0 & 0 & 0 & r \end{bmatrix} \right)$$

Note the default constraint for R is “diagonal and equal” so we can leave this off when specifying the `constraint` argument. To fit, use this code (output not shown):

```
kemfit =
  MARSS(dat, constraint=list(Z=factor(c(1,1,1,1,1))))
```

Success! Parameters converged at 20 iterations.

Alert: `conv.test.slope.tol` is 0.5.

Test with smaller values (<0.1) to ensure convergence.

MARSS fit is

Estimation method: kem

Estimation converged in 20 iterations.

Log-likelihood: 3.575237

AIC: 8.849526    AICc: 11.17211

	Estimate
A.2	0.80614
A.3	0.28714
A.4	-0.54169
A.5	-0.61899
Q.1	0.00458
R.1	0.04512
U.1	0.04763
x0.1	6.38185

Standard errors have not been calculated.  
 Use MARSSparamCIs to compute CIs and bias estimates.

### 6.1.7 One hidden state and five independent observation time series with different variances

Mathematically, this model is:

$$x_t = x_{t-1} + u + v_t, \quad v_t \sim N(0, q)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} x_t + \begin{bmatrix} 0 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix}, \quad \mathbf{w}_t \sim \text{MVN} \left( 0, \begin{bmatrix} r_1 & 0 & 0 & 0 & 0 \\ 0 & r_2 & 0 & 0 & 0 \\ 0 & 0 & r_3 & 0 & 0 \\ 0 & 0 & 0 & r_4 & 0 \\ 0 & 0 & 0 & 0 & r_5 \end{bmatrix} \right)$$

To fit this model:

```
kemfit =
  MARSS(dat, constraint=list(Z=factor(c(1,1,1,1,1)),
    R="diagonal and unequal"))
```

Success! Parameters converged at 21 iterations.  
 Alert: conv.test.slope.tol is 0.5.  
 Test with smaller values (<0.1) to ensure convergence.

MARSS fit is  
 Estimation method: kem  
 Estimation converged in 21 iterations.  
 Log-likelihood: 16.65824  
 AIC: -9.316477    AICc: -3.937167

	Estimate
A.2	0.79435
A.3	0.27418
A.4	-0.53862
A.5	-0.61053
Q.1	0.00607
R.1	0.03217
R.2	0.03506
R.3	0.01361
R.4	0.01095
R.5	0.19662

```
U.1    0.05276
x0.1   6.26683
```

Standard errors have not been calculated.  
Use `MARSSparamCIs` to compute CIs and bias estimates.

### 6.1.8 Two hidden state processes

Here we fit a model with two hidden states (north and south) where observation time series 1-3 are for the north and 4-5 are for the south. We make the hidden state processes independent (meaning a diagonal  $\mathbf{Q}$  matrix) but with the same process variance. We make the observation errors i.i.d. (the default) and the  $\mathbf{u}$  elements equal. Mathematically, this is the model:

$$\begin{bmatrix} x_{n,t} \\ x_{s,t} \end{bmatrix} = \begin{bmatrix} x_{n,t-1} \\ x_{s,t-1} \end{bmatrix} + \begin{bmatrix} u \\ u \end{bmatrix} + \begin{bmatrix} w_{n,t} \\ w_{s,t} \end{bmatrix}, \quad \mathbf{w}_t \sim \text{MVN} \left( 0, \begin{bmatrix} q & 0 \\ 0 & q \end{bmatrix} \right)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_{n,t} \\ x_{s,t} \end{bmatrix} + \begin{bmatrix} 0 \\ a_2 \\ a_3 \\ 0 \\ a_5 \end{bmatrix} + \begin{bmatrix} v_{1,t} \\ v_{2,t} \\ v_{3,t} \\ v_{4,t} \\ v_{5,t} \end{bmatrix}, \quad \mathbf{v}_t \sim \text{MVN} \left( 0, \begin{bmatrix} r & 0 & 0 & 0 & 0 \\ 0 & r & 0 & 0 & 0 \\ 0 & 0 & r & 0 & 0 \\ 0 & 0 & 0 & r & 0 \\ 0 & 0 & 0 & 0 & r \end{bmatrix} \right)$$

To fit the model, use the following code (output not shown):

```
kemfit =
  MARSS(dat, constraint=list(Z=factor(c("N","N","N","S","S")),
    Q="diagonal and equal",U="equal"))
```

## 6.2 Printing and summarizing models and model fits

The package includes print functions for `marssm` objects (model objects) and `marssMLE` objects (fitted models).

```
print(kemfit)
print(kemfit$model)
```

This will print the basic information on model structure and model fit that you have seen in the previous examples.

The package also includes a summary function for models.

```
summary(kemfit$model)
```

Output is not shown because it is verbose, but it prints each matrix with the fixed elements denoted with their values and the free elements denoted by their names. This is very helpful for confirming exactly what model structure you are fitting to the data.

### 6.3 Confidence intervals on a fitted model

The function `MARSSparamCIs()` is used to compute confidence intervals. The function can compute approximate confidence intervals using a numerically estimated Hessian matrix (`method="hessian"`) or via parametric (`method="parametric"`) or non-parametric (`method="innovations"`) bootstrapping.

#### 6.3.1 Approximate confidence intervals from a Hessian matrix

```
#default uses an est Hessian matrix
kem.with.hess.CIs = MARSSparamCIs(kemfit)
```

Use `print` or just type the `marssMLE` object name to see the confidence intervals:

```
print(kem.with.hess.CIs)
```

```
MARSS fit is
Estimation method: kem
Estimation converged in 18 iterations.
Log-likelihood: 7.940938
AIC: 0.1181231   AICc: 2.440704
```

	ML.Est	Std.Err	low.CI	up.CI
A.2	0.79560	0.06146	0.675150	0.9160
A.3	0.27505	0.06247	0.152616	0.3975
A.5	-0.06879	0.08867	-0.242580	0.1050
Q.1	0.00794	0.00435	-0.000578	0.0165
R.1	0.03399	0.00645	0.021351	0.0466
U.1	0.04315	0.01460	0.014536	0.0718
x0.1	6.17126	0.14646	5.884193	6.4583
x0.2	6.20511	0.15805	5.895337	6.5149

CIs calculated at `alpha = 0.05` via `method=hessian`

#### 6.3.2 Confidence intervals from a parametric bootstrap

Use `method="parametric"` to use a parametric bootstrap to compute confidence intervals and bias using a parametric bootstrap.

```
kem.w.boot.CIs=MARSSparamCIs(kemfit,method="parametric",nboot=10)
#nboot should be more like 1000, but set low for example's sake
print(kem.w.boot.CIs)
```

```
MARSS fit is
Estimation method: kem
Estimation converged in 18 iterations.
Log-likelihood: 7.940938
AIC: 0.1181231   AICc: 2.440704
```

	ML.Est	Std.Err	low.CI	up.CI	Est.Bias	Unbias.Est
A.2	0.79560	0.08255	0.6997	0.9157	-0.029975	0.7656
A.3	0.27505	0.06847	0.1982	0.3855	-0.016160	0.2589
A.5	-0.06879	0.08229	-0.1715	0.0796	0.000412	-0.0684
Q.1	0.00794	0.00439	0.0000	0.0122	0.002972	0.0109
R.1	0.03399	0.00783	0.0245	0.0482	0.000849	0.0348
U.1	0.04315	0.01553	0.0228	0.0627	-0.000647	0.0425
x0.1	6.17126	0.14406	5.9234	6.3213	0.039076	6.2103
x0.2	6.20511	0.18298	5.9840	6.5038	-0.057990	6.1471

CI's calculated at  $\alpha = 0.05$  via method=parametric  
 Bias calculated via parametric bootstrapping with 10 bootstraps.

## 6.4 Vectors of just the estimated parameters

Often it is useful to have a vector of the estimated parameters. For example, if you are writing a call to `optim`, you will need a vector of just the estimated parameters.

```
parvec=MARSSvectorizeparam(kemfit)
parvec
```

A.2	A.3	A.5	Q.1	R.1
0.79559952	0.27504984	-0.06878883	0.00794155	0.03398797
U.1	x0.1	x0.2		
0.04315203	6.17125597	6.20511352		

If you want to replace the estimated parameter values with different values, you can use the same function:

```
parvec.new=parvec
parvec.new["Q.1"]=0.02; parvec.new["U.1"]=0.05
kem.new=MARSSvectorizeparam(kemfit, parvec.new)
```

Then you might want to find out the likelihood of the data using those new parameter values. You compute that with the Kalman filter function `MARSSkf()`, sending it the data and the parameters as a list.

```
kf=MARSSkf(dat, kem.new$par, miss.value=-99)
kf$logLik
```

```
[1] 5.943587
```



## 6.5 Degenerate variance estimates

If your data are short relative to the number of parameters you are estimating, then you are liable to find that some of the variance elements are degenerate (equal to zero). Try the following:

```
dat.short = dat[1:4,1:10]
kem.degen = MARSS(dat.short,control=list(allow.degen=FALSE))
```

Warning! Reached maxit before parameters converged. Maxit was 500.

MARSS fit is

Estimation method: kem

WARNING: maxit reached, 500 iter, before convergence.

The likelihood and params are not at the ML values.

Try setting control\$maxit higher.

Log-likelihood: 11.67847

AIC: 2.643054 AICc: 63.30972

	Estimate
Q.1	1.89e-02
Q.2	1.03e-05
Q.3	8.24e-06
Q.4	3.05e-05
R.1	1.22e-02
U.1	9.79e-02
U.2	1.09e-01
U.3	9.28e-02
U.4	1.11e-01
x0.1	5.96e+00
x0.2	6.73e+00
x0.3	6.60e+00
x0.4	5.71e+00

Standard errors have not been calculated.

Use MARSSparamCIs to compute CIs and bias estimates.

Convergence warnings

Warning: the Q.2 parameter value has not converged.

Warning: the Q.3 parameter value has not converged.

Warning: the Q.4 parameter value has not converged.

This will print a warning that the maximum number of iterations was reached before convergence of some of the **Q** parameters. It might be that if you just ran a few more iterations the variances will converge. So first try setting control\$maxit higher.

```
kem.200 = MARSS(dat.short, control=list(maxit=1000,
allow.degen=FALSE), silent=2)
```

Warning! Reached maxit before parameters converged. Maxit was 1000.

Some of the **Q** terms are still not converging. MARSS can detect if a variance is going to zero and it will try zero to see if that has a higher likelihood. Try removing the `allow.degen=FALSE` which was turning off this feature.

```
kem.200 = MARSS(dat.short)
```

Success! Parameters converged at 182 iterations.

Alert: `conv.test.slope.tol` is 0.5.

Test with smaller values (<0.1) to ensure convergence.

MARSS fit is

Estimation method: kem

Estimation converged in 182 iterations.

Log-likelihood: 11.71475

AIC: 2.570491 AICc: 63.23716

	Estimate
Q.1	0.0189
Q.2	0.0000
Q.3	0.0000
Q.4	0.0000
R.1	0.0123
U.1	0.0979
U.2	0.1088
U.3	0.0924
U.4	0.1108
x0.1	5.9645
x0.2	6.7285
x0.3	6.6011
x0.4	5.7093

Standard errors have not been calculated.

Use `MARSSparamCIs` to compute CIs and bias estimates.

So three of the four **Q** elements are going to zero. This often happens when you do not have enough data to estimate both observation and process variance.

Perhaps we are trying to estimate too many variances. We can try using only one variance value in **Q** and one *u* value in **u**:

```
kem.small=MARSS(dat.short,constraint=list(Q="diagonal and equal",
U="equal"))
```

```

Success! Parameters converged at 155 iterations.
Alert: conv.test.slope.tol is 0.5.
Test with smaller values (<0.1) to ensure convergence.

```

```

MARSS fit is
Estimation method: kem
Estimation converged in 155 iterations.
Log-likelihood: 11.18711
AIC: -8.374226   AICc: 0.959107

```

	Estimate
Q.1	0.0000
R.1	0.0194
U.1	0.1030
x0.1	6.0579
x0.2	6.7664
x0.3	6.5293
x0.4	5.7430

```

Standard errors have not been calculated.
Use MARSSparamCIs to compute CIs and bias estimates.

```

No, we still get the convergence warning. There are simply not enough data to estimate both process and observation variances.

## 6.6 Bootstrap parameter estimates

You can easily produce bootstrap parameter estimates from a fitted model using `MARSSboot()`:

```

boot.params = MARSSboot(kemfit,
  nboot=20, output="parameters", sim="parametric")$boot.params

```

```

          |2%      |20%      |40%      |60%      |80%      |100%
Progress: |

```

Use `silent=TRUE` to stop the progress bar from printing. The function will also produce parameter sets generated using a Hessian matrix (`sim="hessian"`) or a non-parametric bootstrap (`sim="innovations"`).

## 6.7 Data simulation

### 6.7.1 Simulated data from a fitted MARSS model

You can easily simulate data from a fitted model using `MARSSsimulate()`.

```
sim.data=MARSSsimulate(kemfit$par, nsim=2, tSteps=100)$sim.data
```

Then you might want to estimate parameters from that simulated data. Above we created two simulated datasets (`nsim=2`). We will fit to the first one. Here the default settings for `MARSS()` are used.

```
kem.sim.1 = MARSS(sim.data[, ,1])
```

Then we might like to see the likelihood of the second set of simulated data under the model fit to the first set of data. We do that with the Kalman filter function.

```
MARSSkf(sim.data[, ,2], kem.sim.1$par, miss.value = -99)$logLik
```

```
[1] 14.56849
```

There are no missing values in our simulated data, but we still need to pass `miss.value` into `MARSSkf()`.

### 6.7.2 Simulated data from a user-built MARSS model

This shows you how to build up a model from scratch and simulate from that using `MARSSsimulate()`.

```
nsim = 20                # number of simulations
burn = 10                # length of burn in period
tSteps = 25
m=3                      #number of hidden state trajectories
B = diag(1,m);
A = array(0, dim=c(m,1))
Z = diag(1,m)
U = array(0.01, dim=c(m,1))
Q = diag(0.3, m)         #independent process errors
R = diag(0.01, m)        #independent obs errors
x0 = array(10, dim=c(m,1)) #initial conditions, really x_1
V0 = array(0,dim=c(m,m))  #leave this 0
the.par.list =
  list(Z=Z, A=A, R=R, B=B, U=U, Q=Q, x0=x0, V0=V0 )
# simulate data
sim = MARSSsimulate(the.par.list,nsim=nsim,tSteps=burn+tSteps)
# take off the burn
obs = sim$sim.data[, (burn+1):(burn+tSteps),] #m x T x nsim
```

### 6.7.3 Correlation between estimated parameters

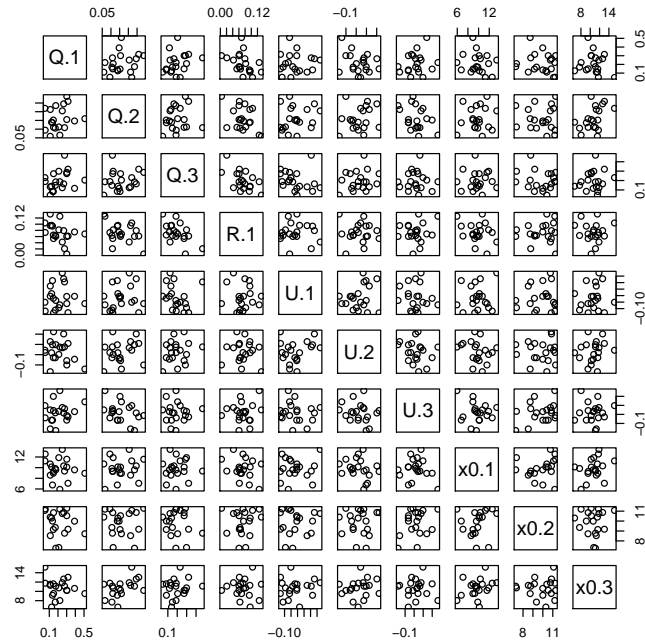
We can use a for loop along with `MARSSvectorizeparam()` to estimate parameters for each of the `nsim` simulated datasets in `obs`, and then assemble the estimates into a matrix.

```

for(i in 1:nsim){
  dat=obs[,i]
  kem.sim=MARSS(dat, silent=TRUE)
  if(i==1) par.sim=MARSSvectorizeparam(kem.sim)
  else par.sim=rbind(par.sim, MARSSvectorizeparam(kem.sim))
}

```

Then we use `pairs()` to get a quick visual look at how the parameters are correlated (or not) and how variable they are (Figure 6.1).



**Fig. 6.1.** The parameter pairs plotted against each other using the command `pairs(par.sim)`.

We could also use `MARSSboot()` to do this. However, `MARSSboot()` is designed to take a `marssMLE` object (such as would be returned from `MARSS()`).

## 6.8 Bootstrap AIC

The function `MARSSaic()` computes a bootstrap AIC for model selection purposes. Use `output="AICbp"` to produce a parameter bootstrap. Use `output="AICbb"` to produce a non-parameter bootstrap AIC.

```
kemfit.with.AICb = MARSSaic(kemfit, output = "AICbp",
  Options = list(nboot = 10, silent=TRUE))
#nboot should be more like 1000, but set low here for example sake
```

```
print(kemfit.with.AICb)
```

MARSS fit is

Estimation method: kem

Estimation converged in 18 iterations.

Log-likelihood: 7.940938

AIC: 0.1181231    AICc: 2.440704    AICbp(param): 6.94647

	Estimate
A.2	0.79560
A.3	0.27505
A.5	-0.06879
Q.1	0.00794
R.1	0.03399
U.1	0.04315
x0.1	6.17126
x0.2	6.20511

Standard errors have not been calculated.

Use `MARSSparamCIs` to compute CIs and bias estimates.

## Case study instructions

The case studies walk you through some analyses of multivariate population count data using MARSS models and the `MARSS()` function. This will take you through both the conceptual steps (with pencil and paper) and a *R* step which translates the conceptual model into code.

### Set-up

- If you haven't already, install the MARSS package. See directions on the CRAN webpage (<http://cran.r-project.org/>) for instructions on installing packages. You will need write permissions for your *R* program directories to install packages. See the help pages on CRAN for workarounds if you don't have write permission.
- Type in `library(MARSS)` at the *R* command line. Now you should be ready.
- Each case study comes with an associated script file. To open up a copy of the case study script with the code you need to do the exercises, type `show.doc(MARSS, Case_study_#.R)` (with `#` replaced by the case study number).

### Tips

- `summary(foo$model)`, where `foo` is a fitted model object, will print detailed information on the structure of the MARSS model that was fit in the call `foo = MARSS(logdata)`. This allows you to double check the model you fit. `print(foo)` will print a 'English' version of the model structure along with the parameter estimates.
- When you run `MARSS()`, it will output the number of iterations used. If you reached the maximum, re-run with `control=list(maxit=...)` set higher than the default.

- If you mis-specify the model, `MARSS()` will post an error that should give you an idea of the problem (make sure `silent=FALSE` to see full error reports). Remember, the number of rows in your data is  $n$ , time is across the columns, and the length of the vector or factors passed in for `constraint$Z` must be  $m$ , the number of  $x$  hidden state trajectories in your model.
- If you are fitting to population counts, your data must be logged (base e) before being passed in. The default missing value indicator is -99. You can change that by passing in `miss.value=...`.
- Running `MARSS(data)`, with no arguments except your data, will fit a MARSS model with  $m = n$ , a diagonal **Q** matrix with  $m$  variances, and i.i.d. observation errors.



## Case Study 1: Count-based Population Viability Analysis using corrupted data

### 8.1 The Problem

Estimates of extinction and quasi-extinction risk are an important risk metric used in the management and conservation of endangered and threatened species. By necessity, these estimates are based on data that contain both variability due to real year-to-year changes in the population growth rate (process errors) and variability in the relationship between the true population size and the actual count (observation errors). Classic approaches to extinction risk assume the data have only process error, i.e. no observation error. In reality, observation error is ubiquitous both because of the sampling variability and also because of year-to-year (and day-to-day) variability in sightability.

In this case study, we use the Kalman filter to fit a univariate (meaning one time series) state-space model to count data for a population. We will compute the extinction risk metrics given in Dennis et al. (1991), however instead of using a process-error only model (as is done in the original paper), we use a model with both process and observation error. The risk metrics and their interpretations are the same as in Dennis et al. (1991). The only real difference is how we compute  $\sigma^2$ , the process error variance. However this difference has a large effect on our risk estimates, as you will see.

In this case study, we use a density-independent model, which is the same as the Gompertz model (Equation 3.1) with  $\mathbf{B} = 1$ . Density-independence is often a reasonable assumption when doing a population viability analysis because we do such calculations for at-risk populations that are either declining or that are well below historical levels (and presumably carrying capacity). In an actual population viability analysis, it is necessary to justify this assumption and if there is reason to doubt the assumption, one tests for density-dependence (Taper and Dennis, 1994) and does sensitivity analyses using state-space models with density-dependence (Dennis et al., 2006).

The univariate model is written:

$$x_t = x_{t-1} + u + w_t \quad \text{where } w_t \sim N(0, \sigma^2) \quad (8.1)$$

$$y_t = x_t + v_t \quad \text{where } v_t \sim N(0, \eta^2) \quad (8.2)$$

where  $y_t$  is the logarithm of the observed population size at time  $t$ ,  $x_t$  is the unobserved state at time  $t$ ,  $u$  is the growth rate, and  $\sigma^2$  and  $\eta^2$  are the process and observation error variances, respectively. In the *R* code to follow,  $\sigma^2$  is denoted **Q** and  $\eta^2$  is denoted **R** because the functions we are using are also for multivariate state-space models and those models use **Q** and **R** for the respective variance-covariance matrices.

## 8.2 Simulated data with process and observation error

We will start by using simulated data to see the difference between data and estimates from a model with process error only versus a model that also includes observation error. For our simulated data, we used a decline of 5% per year, process variability of 0.01 (typical for big mammals), and a observation variability of 0.05 (which is a bit on the high end). We'll randomly set 10% of the values as missing. Here is the code:

First, set things up:

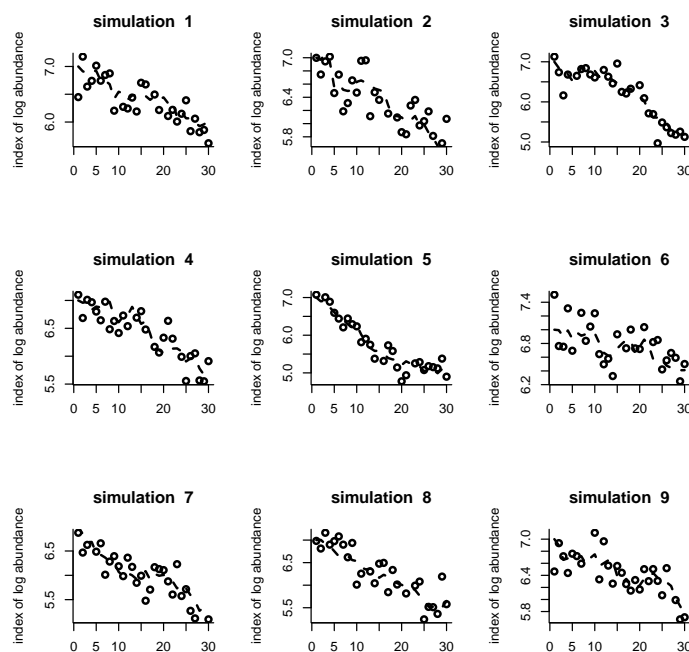
```
sim.u = -0.05      # growth rate
sim.Q = 0.01       # process error variance
sim.R = 0.05       # non-process error variance
nYr= 30            # number of years of data to generate
fracmissing = 0.1  # fraction of years that are missing
init = 7           # log of initial pop abundance
years = seq(1:nYr) # sequence 1 to nYr
x = rep(NA,nYr)    # replicate NA nYr times
y = rep(NA,nYr)
```

Then generate the population sizes using Equation 8.1:

```
x[1]=init
for(t in 2:nYr){
  x[t] = x[t-1]+ sim.u + rnorm(1,mean=0,sd=sqrt(sim.Q)) }
```

Lastly, add observation error using Equation 8.2 and then add missing values:

```
for(t in 1:nYr){
  y[t]= x[t] + rnorm(1,mean=0,sd=sqrt(sim.R))
}
missYears = sample(years[2:(nYr-1)],floor(fracmissing*nYr),
  replace = FALSE)
y[missYears]=-99
```



**Fig. 8.1.** Plot of nine simulated population time series with process and observation error. Circles are observation and the dashed line is the true population size.

Stochastic population trajectories show much variation, so it is best to look at a few simulated data sets at once. In Figure 8.1, nine simulations from the identical parameters are shown.

### Example 8.1 (The effect of parameter values on parameter estimates)

*A good way to get a feel for reasonable  $\sigma^2$  values is to generate simulated data and look at the time series. As a biologist, you probably have a pretty good idea of what kind of year-to-year population changes are reasonable for your species. For example for many large mammalian species, the maximum population yearly increase would be around 50% (the population could go from 1000 to 1500 in one year), but some of fish species could easily double or even triple in a really good year. Your observed data may bounce around a lot for many different reasons having to do with sightability, sampling error, age-structure, etc., but the underlying population trajectory is constrained by the*

kinds of year-to-year changes in population size that are biologically possible for your species.  $\sigma^2$  describes those true population changes.

Run the exercise code several times using different parameter values to get a feel for how different the time series can look based on identical parameter values. You can cut and paste from the pdf into the R command line. Typical vertebrate  $\sigma^2$  values are 0.002 to 0.02, and typical  $\eta^2$  values are 0.005 to 0.1. A  $u$  of -0.01 translates to an average 1% per year decline and a  $u$  of -0.1 translates to an average 10% per year decline (approximately).

---

### Example 8.1 code

Type `show.doc(MARSS, Case_study_1.R)` to open a file with all the example code.

```
par(mfrow=c(3,3))
sim.u = -0.05
sim.Q = 0.01
sim.R = 0.05
nYr= 30
fracmiss = 0.1
init = 7
years = seq(1:nYr)
for(i in 1:9){
  x = rep(NA,nYr) # vector for ts w/o measurement error
  y = rep(NA,nYr) # vector for ts w/ measurement error
  x[1]=init
  for(t in 2:nYr){
    x[t] = x[t-1]+ sim.u + rnorm(1, mean=0, sd=sqrt(sim.Q)) }
  for(t in 1:nYr){
    y[t]= x[t] + rnorm(1,mean=0,sd=sqrt(sim.R)) }
  missYears =
    sample(years[2:(nYr-1)],floor(fracmiss*nYr),replace = FALSE)
  y[missYears]=-99
  plot(years[y!=-99], y[y!=-99],
        xlab="",ylab="log abundance",lwd=2,bty="l")
  lines(years,x,type="l",lwd=2,lty=2)
  title(paste("simulation ",i) )
}
legend("topright", c("Observed","True"),
      lty = c(-1, 2), pch = c(1, -1))
```

---

## 8.3 Maximum-likelihood parameter estimation

### 8.3.1 Model with process and observation error

We put the simulated data through the Kalman-EM algorithm in order to estimate the parameters,  $u$ ,  $\sigma^2$ , and  $\eta^2$ , and population sizes. These are the estimates using a model with process and observation variability. The function call is `kem = MARSS(data)`, where `data` is a vector of logged (base e) counts with missing values denoted by -99. After this call, the maximum-likelihood parameter estimates are `kem$par$U`, `kem$par$Q` and `kem$par$R`. There are numerous other outputs from the `MARSS()` function. To get a list of the outputs type in `names(kem)`. Note that `kem` is just a name; the output could have been called `foo`. Here's code to fit to the simulated time series:

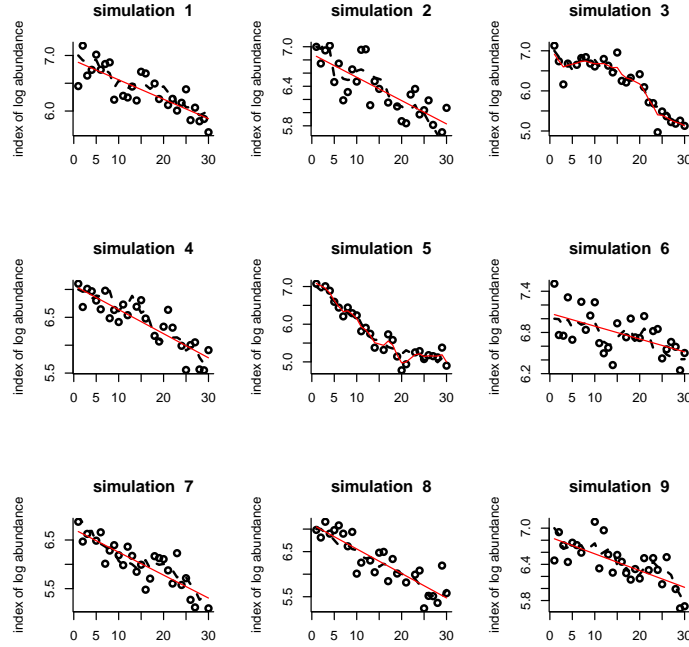
```
kem = MARSS(y)
```

Let's look at the parameter estimates for the nine simulated time series in Figure 8.1 to get a feel for the variation. The `MARSS()` function was used on each time series to produce parameter estimate for each simulation. The estimates are followed by the mean (over the nine simulations) and the true values:

	kem.U	kem.Q	kem.R
sim 1	-0.03494735	0.000000000	0.04984165
sim 2	-0.03539361	0.000000000	0.05702256
sim 3	-0.06167768	0.019339366	0.04149535
sim 4	-0.04325879	0.000000000	0.04575033
sim 5	-0.07187786	0.022111208	0.01807441
sim 6	-0.01871904	0.000000000	0.06201101
sim 7	-0.04714729	0.000000000	0.05943489
sim 8	-0.05436225	0.000000000	0.07322492
sim 9	-0.02775879	0.000000000	0.05468758
mean sim	-0.04390474	0.004605619	0.05128252
true	-0.05000000	0.01000000	0.05000000

As expected, the estimated parameters do not exactly match the true parameters, but the average should be fairly close (although nine simulations is a small sample size). Also note that although we do not get  $u$  quite right, our estimates are usually negative. Thus our estimates usually indicate declining dynamics.

The Kalman-EM algorithm also gives an estimate of the true population size with observation error removed. This is in `kem$states`. Figure 8.2 shows the estimated true states of the population over time as a solid line. Note that the solid line is considerably closer to the actual true states (dashed line) than the observations. On the other hand with certain datasets, the estimates can be quite wrong as well!



**Fig. 8.2.** The circles are the observed population sizes with error. The dashed lines are the true population sizes. The solid thin lines are the estimates of the true population size from the Kalman-EM algorithm.

### 8.3.2 Model with no observation error

We used the Kalman-EM algorithm to estimate the mean population rate  $u$  and process variability  $\sigma^2$  under the assumption that the count data have observation error. However, the classic approach to this problem, referred to as the “Dennis model” (Dennis et al., 1991), uses a model that assumes the data have no observation error; all the variability in the data is assumed to result from process error. This approach works well if the observation error in the data is low, but not so well if the observation error is high. We will next fit the data using the classic approach so that we can compare and contrast parameter estimates from the different methods.

Using the estimation method in Dennis et al. (1991), our data need to be re-specified as the observed population changes (`delta.pop`) between censuses along with the time between censuses (`tau`). We re-specify the data as follows:

```
den.years = years[y!=-99] # the non missing years
den.y = y[y!=-99] # the non missing counts
den.n.y = length(den.years)
```

```

delta.pop = rep(NA, den.n.y-1 ) # population transitions
tau = rep(NA, den.n.y-1 ) # step sizes
for (i in 2:den.n.y ){
  delta.pop[i-1] = den.y[i] - den.y[i-1]
  tau[i-1] = den.years[i] - den.years[i-1]
} # end i loop


```

Next, we regress the changes in population size between censuses (`delta.pop`) on the time between censuses (`tau`) while setting the regression intercept to 0. The slope of the resulting regression line is an estimate of  $u$ , while the variance of the residuals around the line is an estimate of  $\sigma^2$ . The regression is shown in Figure 8.3. Here is the code to do that regression:

```

den91 <- lm(delta.pop ~ -1 + tau)
# note: the "-1" specifies no intercept
den91.u = den91$coefficients
den91.Q = var(resid(den91))

```



**Fig. 8.3.** The regression of  $\log(N_{t+\tau}) - \log(N_t)$  against  $\tau$ . The slope is the estimate of  $u$  and the variance of the residuals is the estimate of  $\sigma^2$ .

Here are the parameter values for the data in Figure 8.2 using the process-error only model:

	den91.U	den91.Q
sim 1	-0.029996777	0.10075189
sim 2	-0.037233652	0.11118126
sim 3	-0.050171644	0.10713363
sim 4	-0.047877065	0.08761822
sim 5	-0.056771766	0.05968257
sim 6	-0.025068284	0.13324756
sim 7	-0.052921562	0.10977084
sim 8	-0.039419954	0.17085967
sim 9	-0.007105354	0.12140355
mean sim	-0.038507340	0.11129436
true	-0.050000000	0.01000000

Notice that the  $u$  estimates are similar to those from the Kalman-EM algorithm, but the  $\sigma^2$  estimate ( $Q$ ) is much larger. That is because this approach treats all the variance as process variance, so any observation variance in the data is lumped into process variance (in fact it appears as an additional variance of twice the observation variance).

### Example 8.2 (The variability in parameter estimates)

*In this example, you will look at how variable the parameter estimates are by generating multiple simulated data sets and then estimating parameter values for each. You'll compare the Kalman-EM estimates to the estimates using a process error only model (i.e. ignoring the observation error).*



**Example 8.2 code**

Type `show.doc(MARSS, Case_study_1.R)` to open a file with all the example code.

You will not be able to edit this file. To edit, copy and paste into a new script file.

---

```

sim.u = -0.05 # growth rate
sim.Q = 0.01  # process error variance
sim.R = 0.05  # non-process error variance
nYr= 30       # number of years of data to generate
fracmiss = 0.0 # fraction of years that are missing
init = 7      # log of initial pop abundance (~1100 individuals)
nsim = 9
years = seq(1:nYr) # col of years
params = matrix(NA, nrow=(nsim+2), ncol=5,
  dimnames=list(c(paste("sim",1:nsim),"mean sim","true"),
c("kem.U","den91.U","kem.Q","kem.R", "den91.Q")))
x.ts = matrix(NA,nrow=nsim,ncol=nYr) # ts w/o measurement error
y.ts = matrix(NA,nrow=nsim,ncol=nYr) # ts w/ measurement error
for(i in 1:nsim){
  x.ts[i,1]=init
  for(t in 2:nYr){
    x.ts[i,t] = x.ts[i,t-1]+sim.u+rnorm(1,mean=0,sd=sqrt(sim.Q))}
  for(t in 1:nYr){
    y.ts[i,t] = x.ts[i,t]+rnorm(1,mean=0,sd=sqrt(sim.R))}
  missYears = sample(years[2:(nYr-1)], floor(fracmiss*nYr),
    replace = FALSE)
  y.ts[i,missYears]==-99

  #Kalman-EM estimates
  #So that this code runs fast, we're using a rough convergence test (abstol)
  #change that line to kem = MARSS(y.ts[i,], silent=TRUE)
  #for a better (albeit slower) convergence test
  kem = MARSS(y.ts[i,], control=list(abstol=0.01), silent=TRUE)
  params[i,c(1,3,4)] = c(kem$par$U,kem$par$Q,kem$par$R)

  #Dennis et al 1991 estimates
  den.years = years[y.ts[i,]!=-99] # the non missing years
  den.yts = y.ts[i,y.ts[i,]!=-99] # the non missing counts
  den.n.yts = length(den.years)
  delta.pop = rep(NA, den.n.yts-1 ) # transitions
  tau = rep(NA, den.n.yts-1 ) # time step lengths
  for (t in 2:den.n.yts ){
    delta.pop[t-1] = den.yts[t] - den.yts[t-1] # transitions
    tau[t-1] = den.years[t]-den.years[t-1] # time step length
  } # end i loop
  den91 <- lm(delta.pop ~ -1 + tau) # -1 specifies no intercept
  params[i,c(2,5)] = c(den91$coefficients, var(resid(den91)))
}
params[nsim+1,]=apply(params[1:nsim,],2,mean)
params[nsim+2,]=c(sim.u,sim.u,sim.Q,sim.R,sim.Q)

```

---

Here is an example of the output from the code:

```
print(params,digits=3)

      kem.U den91.U      kem.Q kem.R den91.Q
sim 1  -0.0427 -0.0493 0.000253 0.0454 0.1054
sim 2  -0.0784 -0.0731 0.000159 0.0417 0.1134
sim 3  -0.0442 -0.0536 0.006801 0.0578 0.1394
sim 4  -0.0607 -0.0516 0.004877 0.0919 0.1781
sim 5  -0.0666 -0.0658 0.000313 0.0391 0.0664
sim 6  -0.0587 -0.0451 0.009608 0.0338 0.0766
sim 7  -0.0216 -0.0325 0.000345 0.0497 0.1075
sim 8  -0.0589 -0.0531 0.001452 0.0647 0.1502
sim 9  -0.0558 -0.0604 0.012751 0.0496 0.1269
mean sim -0.0542 -0.0538 0.004062 0.0526 0.1182
true    -0.0500 -0.0500 0.010000 0.0500 0.0100
```

1. Re-run the code a few times to see the performance of the estimates using a state-space model (`kem`) versus the model with no observation error (`den91`). You can copy and paste the code from the pdf file into R .
2. Alter the observation variance, `sim.R`, in the data generation step in order to get a feel for performance as observations are further corrupted. What happens as observation error is increased?
3. Decrease the number of years of data, `nYr`, and re-run the parameter estimation. What changes?

If you find that the exercise code takes too long to run, reduce the number of simulations (by reducing `nsim` in the code).

## 8.4 Probability of hitting a threshold $\Pi(x_d, t_e)$

A common extinction risk metric is ‘the probability that a population will hit a certain threshold  $x_d$  within a certain time frame  $t_e$  – if the observed trends continue’. In practice, the threshold used is not  $N_e = 1$ , which would be true extinction. Often a ‘functional’ extinction threshold will be used ( $N_e \gg 1$ ). Other times a threshold representing some fraction of current levels is used. The latter is used because we often have imprecise information about the relationship between the true population size and what we measure in the field; that is, many population counts are index counts. In these cases, one

must use ‘fractional declines’ as the threshold. Also, extinction estimates that use an absolute threshold (like 100 individuals) are quite sensitive to error in the estimate of true population size. Here, we are going to use fractional declines as the threshold, specifically  $p_d = 0.1$  which means a 90% decline.

The probability of hitting a threshold, denoted  $\Pi(x_d, t_e)$ , is typically presented as a curve showing the probabilities of hitting the threshold ( $y$ -axis) over different time horizons ( $t_e$ ) on the  $x$ -axis. Extinction probabilities can be computed through Monte Carlo simulations or analytically using Equation 16 in Dennis et al. (1991) (note there is a typo in Equation 16; the last  $+$  is supposed to be a  $-$ ). We will use the latter method:

$$\Pi(x_d, t_e) = \pi(u) \times \Phi\left(\frac{-x_d + |u|t_e}{\sqrt{\sigma^2 t_e}}\right) + \exp(2x_d|u|/\sigma^2) \Phi\left(\frac{-x_d - |u|t_e}{\sqrt{\sigma^2 t_e}}\right) \quad (8.3)$$

where  $x_e$  is the threshold and is defined as  $x_e = \log(N_0/N_e)$ .  $N_0$  is the current population estimate and  $N_e$  is the threshold. If we are using fractional declines then  $x_e = \log(N_0/(p_d \times N_0)) = -\log(p_d)$ .  $\pi(u)$  is the probability that the threshold is eventually hit (by  $t_e = \infty$ ).  $\pi(u) = 1$  if  $u \leq 0$  and  $\pi(u) = \exp(-2ux_d/\sigma^2)$  if  $u > 0$ .  $\Phi()$  is the cumulative probability distribution of the standard normal (mean = 0, sd = 1).

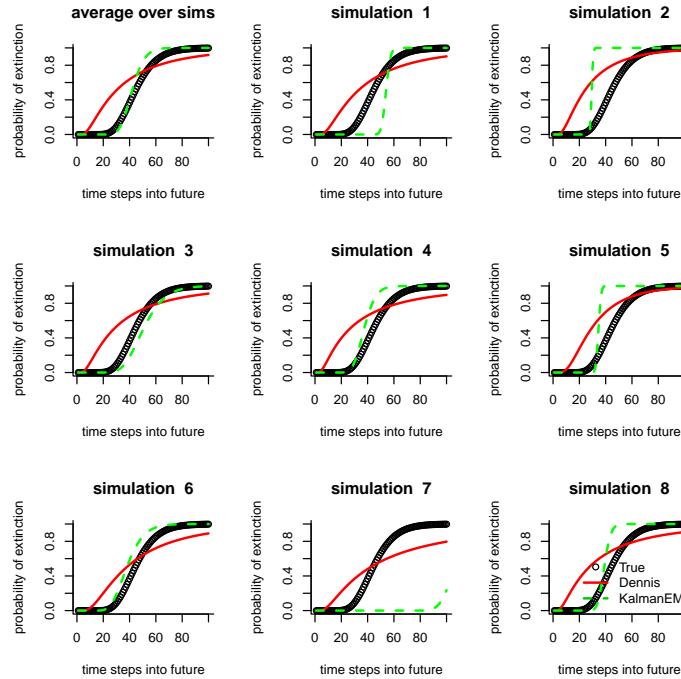
Here is the *R* code for that computation:

```
pd = 0.1 #means a 90 percent decline
tyrs = 1:100
xd = -log(pd)
p.ever = ifelse(u<=0,1,exp(-2*u*xd/Q)) #Q=sigma2
for (i in 1:100){
  Pi[i] = p.ever * pnorm((-xd+abs(u)*tyrs[i])/sqrt(Q*tyrs[i]))+
    exp(2*xd*abs(u)/Q)*pnorm((-xd-abs(u)*tyrs[i])/sqrt(Q*tyrs[i]))
}
```

Figure 8.4 shows the estimated probabilities of hitting the 90% decline for the nine 30-year times series simulated with  $u = -0.05$ ,  $\sigma^2 = 0.01$  and  $\eta^2 = 0.05$ . The dashed line shows the estimates using the Kalman-EM parameter estimates and the solid line shows the estimates using a process-error only model (the **den91** estimates). The circles are the true probabilities. The difference between the estimates and the true probabilities is due to errors in  $\hat{u}$ . Those errors are due largely to process error—not observation error. As we saw earlier, by chance population trajectories with a  $u < 0$  will increase, even over a 30-year period. In this case,  $\hat{u}$  will be positive when in fact  $u < 0$ .

Looking at the figure, it is obvious that the probability estimates are highly variable. However, look at the first panel. This is the average estimate (over nine simulations). Note that on average (over nine simulations), the estimates are good. If we had averaged over 1000 simulations instead of nine, you would see that the Kalman-EM line falls on the true line. It is an unbiased predictor. While that may seem a small consolation if estimates for individual simulations are all over the map, it is important for correctly specifying our uncertainty

about our estimates. Second, rather than focusing on how the estimates and true lines match up, see if there are any types of forecasts that seem better than others. For example, are 20-year predictions better than 50-year and are 100-year forecasts better or worse. In Exercise 3, you will remake this figure with different  $u$ . You'll discover from that forecasts are more certain for populations that are declining faster.



**Fig. 8.4.** Plot of the true and estimated probability of declining 90% in different time horizons for nine simulated population time series with observation error. The plot may look like a step-function if the  $\sigma^2$  estimate is very small ( $<1e-4$  or so).

### Example 8.3 (The effect of parameter values on risk estimates)

*In this example, you will recreate Figure 8.4 using different parameter values. This will give you a feel for how variability in the data and population process affect the risk estimates. You'll need to run the Example 8.2 code before running the Example 8.3 code.*

**Example 8.3 code**

Type `show.doc(MARSS, Case_study_1.R)` to open a file with all the example code.

---

```
#Needs Exercise 2 to be run first
par(mfrow=c(3,3))
pd = 0.1; xd = -log(pd) # decline threshold
te = 100; tyrs = 1:te # extinction time horizon
for(j in c(10,1:8)){
  real.ex = denn.ex = kal.ex = matrix(nrow=te)

  #Kalman-EM parameter estimates
  u=params[j,1]; Q=params[j,3]
  p.ever = ifelse(u<=0,1,exp(-2*u*xd/Q))
  for (i in 1:100){
    if(is.finite(exp(2*xd*abs(u)/Q))){
      sec.part = exp(2*xd*abs(u)/Q)*pnorm((-xd-abs(u)* tyrs[i])/sqrt(Q*tyrs[i]))
    }else sec.part=0
    kal.ex[i]=p.ever*pnorm((-xd+abs(u)*tyrs[i])/sqrt(Q*tyrs[i]))+sec.part
  } # end i loop

  #Dennis et al 1991 parameter estimates
  u=params[j,2]; Q=params[j,5]
  p.ever = ifelse(u<=0,1,exp(-2*u*xd/Q))
  for (i in 1:100){
    denn.ex[i]=p.ever*pnorm((-xd+abs(u)*tyrs[i])/sqrt(Q*tyrs[i]))+
      exp(2*xd*abs(u)/Q)*pnorm((-xd-abs(u)*tyrs[i])/sqrt(Q*tyrs[i]))
  } # end i loop

  #True parameter values
  u=sim.u; Q=sim.Q
  p.ever = ifelse(u<=0,1,exp(-2*u*xd/Q))
  for (i in 1:100){
    real.ex[i]=p.ever*pnorm((-xd+abs(u)*tyrs[i])/sqrt(Q*tyrs[i]))+
      exp(2*xd*abs(u)/Q)*pnorm((-xd-abs(u)*tyrs[i])/sqrt(Q*tyrs[i]))
  } # end i loop

  #plot it
  plot(tyrs, real.ex, xlab="time steps into future",
    ylab="probability of extinction", ylim=c(0,1), bty="l")
  if(j<=8) title(paste("simulation ",j) )
  if(j==10) title("average over sims")
  lines(tyrs,denn.ex,type="l",col="red",lwd=2,lty=1)
  lines(tyrs,kal.ex,type="l",col="green",lwd=2,lty=2)
}
legend("bottomright",c("True","Dennis","KalmanEM"),pch=c(1,-1,-1),
  col=c(1,2,3),lty=c(-1,1,2),lwd=c(-1,2,2),bty="n")
```

---


1. Change `sim.R` and rerun the Example 8.2 code. Then run the Example 8.3 code. When are the estimates using the process-error only model (`den91`) worse and in what way are they worse?
2. You might imagine that you should always use a model that includes observation error, since in practice observations are never perfect. However, there is a cost to estimating that extra variance parameter and the cost is a more variable  $\sigma^2$  (Q) estimate. Play with shortening the time series and decreasing the `sim.R` values. Are there situations when the ‘cost’ of the extra parameter is greater than the ‘cost’ of ignoring observation error?
3. How does changing the extinction threshold (`pd`) change the extinction probability curves? (Do not remake the data, i.e. don’t rerun the Example 8.2 code.)
4. How does changing the rate of decline (`sim.u`) change the estimates of risk? Rerun the Example 8.2 code using a lower `u`; this will create a new matrix of parameter estimates. Then run the Example 8.3 code. Do the estimates seem better or worse for rapidly declining populations?
5. Rerun the Example 8.2 code using fewer number of years (`nYr` smaller) and increase `fracmiss`. Then run the Example 8.3 code. The graphs will start to look peculiar. Why do you think it is doing that? Hint: look at the estimated parameters.

## 8.5 Certain and uncertain regions

From Example 8.3, you have observed one of the problems with estimates of the probability of hitting thresholds. Looking over the nine simulations, your risk estimates will be on the true line sometimes and other times they are way off. So your estimates are variable and one should not present only the point estimates of the probability of 90% decline. At the minimum, confidence intervals need to be added (next section), but even with confidence intervals, the probability of hitting declines often does not capture our certainty and uncertainty about extinction risk estimates.

From Example 8.3, you might have also noticed that there are some time horizons (10, 20 years) for which the estimate are highly certain (the threshold is never hit), while for other time horizons (30, 50 years) the estimates are all over the map. Put another way, you may be able to say with high confidence that a 90% decline will not occur between years 1 to 20 and that by year 100 it most surely will have occurred. However, between the years 20 and 100, you are very uncertain about the risk. The point is that you can be certain about some forecasts while at the same time being uncertain about other forecasts.

One way to show this is to plot the uncertainty as a function of the forecast, where the forecast is defined in terms of the forecast length (number of years) and forecasted decline (percentage). Uncertainty is defined as how much of the 0-1 range your 95% confidence interval covers. Ellner and Holmes (2008) show such a figure (their Figure 1). Figure 8.5 shows a version of this figure that you can produce with the function `CSEGtmfigure(u= val, N= val, s2p= val)`. For the figure, the values  $u = -0.05$  which is a 5% per year decline,  $N = 25$  so 25 years between the first and last census, and  $s_p^2 = 0.01$  are used. The process variability for big mammals is typically in the range of 0.002 to 0.02.



**Fig. 8.5.** This figure shows your region of high uncertainty (dark grey). In this region, the minimum 95% confidence intervals (meaning if you had no observation error) span 80% of the 0 to 1 probability. That is, you are uncertain if the probability of a specified decline is close to 0 or close to 1. The green (dots) shows where your upper 95% CIs does not exceed  $P=0.05$ . So you are quite sure the probability of a specified decline is less than 0.05. The red (dots) shows where your lower 95% confidence interval is above  $P=.95$ . So you are quite sure the probability is greater than  $P=0.95$ . The light grey is between these two certain/uncertain extremes.

**Example 8.4 (Uncertain and certain regions)**

Use the Example 8.4 code to re-create Figure 8.5 and get a feel for when risk estimates are more certain and when they are less certain.  $N$  are the number of years of data,  $u$  is the mean population growth rate, and  $s2p$  is the process variance.

**Exercise 8.4 code**

Type `show.doc(MARSS, Case_study_1.R)` to open a file with all the example code.

---

```
par(mfrow = c(1, 1))
CSEgtmfigure(N = 30, u = -0.05, s2p = 0.01)
```

---

**8.6 More risk metrics and some real data**

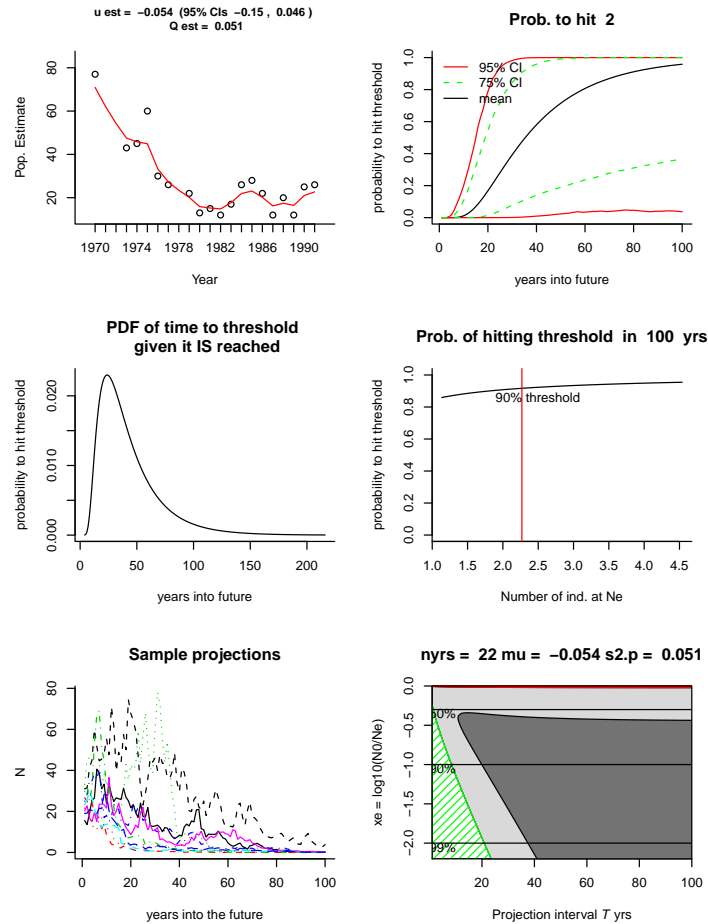
The previous sections have focused on the probability of hitting thresholds because this is an important and common risk metric used in population viability analysis and it appears in IUCN Red List criteria. However, as you have seen, there is high uncertainty associated with such estimates. Part of the problem is that probability is constrained to be 0 to 1, and it is easy to get estimates with confidence intervals that span 0 to 1. Other metrics of risk,  $\hat{u}$  and the distribution of the time to hit a threshold (Dennis et al., 1991), do not have this problem and may be more informative. Figure 8.6 shows different risk metrics from Dennis et al. (1991) on a single plot. This figure is generated by a call to the function `CSEGriskfigure()`:

```
dat=read.table(datafile, skip=1)
dat=as.matrix(dat)
CSEGriskfigure(dat)
```

The `datafile` is the name of the data file, with years in column 1 and population count (logged) in column 2. `CSEGriskfigure()` has a number of arguments that can be passed in to change the default behavior. The variable `te` is the forecast length (default is 100 years), `threshold` is the extinction threshold either as an absolute number, if `absolutethresh=TRUE`, or as a fraction of current population count, if `absolutethresh=FALSE`. The default is `absolutethresh=FALSE` and `threshold=0.1`. `datalogged=TRUE` means the data are already logged; this is the default.

**Example 8.5 (Risk figures for different species)**





**Fig. 8.6.** Risk figure using data for the critically endangered African Wild Dog (data from Ginsberg et al. 1995). This population went extinct after 1992.

Use the Example 8.5 code to re-create Figure 8.6. The package includes other data for you to run: `prairiechicken` from the endangered Attwater Prairie Chicken, `graywhales` from Gerber et al. (1999), and `grouse` from the Sharp-tailed Grouse (a species of U.S. federal concern) in Washington State. Note for some of these other datasets, the Hessian matrix cannot be inverted and you will need to use `CI.method="parametric"`. If you have other text files of data, you can run those too. The commented lines show how to read in data from a tab-delimited text file with a header line.

---

**Exercise 5 code**

Type `show.doc(MARSS, Case_study_1.R)` to open a file with the example code.

```
#If you have your data in a tab delimited file with a header
#This is how you would read it in using file.choose()
#to call up a directory browser.
#However, the package has the datasets for the exercises
#dat=read.table(file.choose(), skip=1)
#dat=as.matrix(dat)
dat = wilddogs
CSEGriskfigure(dat, CI.method="hessian", silent=TRUE)
```

---

## 8.7 Confidence intervals

The figures produced by `CSEGriskfigure()` have confidence intervals (95% and 75%) on the probabilities in the top right panel. A standard way to produce these intervals is via parametric bootstrapping. Here are the steps in a parametric bootstrap:

- You estimate  $u$ ,  $\sigma^2$  and  $\eta^2$
- Then you simulate time series using those estimates and Equations 8.1 and 8.2
- Then you re-estimate your parameters from the simulated data (using say `MARSS(simdata)`)
- Repeat for 1000s of time series simulated using your estimated parameters. This gives you a large set of bootstrapped parameter estimates
- For each bootstrapped parameter set, compute a set of extinction estimates (you use Equation 8.3 and code from Example 8.3)
- The  $\alpha\%$  ranges on those bootstrapped extinction estimates gives you your  $\alpha$  confidence intervals on your probabilities of hitting thresholds

The MARSS package provides the function `MARSSparamCIs()` to add bootstrapped confidence intervals to fitted models (type `?MARSSparamCIs` to learn about the function).

In the function `CSEGriskfigure()`, you can set `CI.method = c("hessian", "parametric", "innovations", "none")` to tell it how to compute the confidence intervals. The methods ‘parametric’ and ‘innovations’ specify parametric and non-parametric bootstrapping respectively. Producing parameter estimates by bootstrapping is quite slow. Approximate confidence intervals on the parameters can be generated rapidly using the inverse of a numerically estimated Hessian matrix (method ‘hessian’). This uses an estimate of the

variance-covariance matrix of the parameters (the inverse of the Hessian matrix). Using an estimated Hessian matrix to compute confidence intervals is a handy trick that can be used for all sorts of maximum-likelihood parameter estimates.

## 8.8 Comments

Data with cycles, from age-structure or predator-prey interactions, are difficult to analyze and the Kalman-EM algorithm used in the MARSS package will give poor estimates for this type of data. The slope method (Holmes, 2001) is robust to those problems. Holmes et al. (2007) used the slope method in a large study of data from endangered and threatened species, and Ellner and Holmes (2008) showed that the slope estimates are close to the theoretical minimum uncertainty. Especially, when doing a population viability analysis using a time series with fewer than 25 years of data, the slope method is often less biased and (much) less variable because that method is less data-hungry (Holmes, 2004). However the slope method is not a true maximum-likelihood method and thus constrains the types of further analyses you can do (such as model selection).



## Case study 2: Combining multi-site and subpopulation data to estimate regional population dynamics

### 9.1 The problem

In this case study, we will use multivariate state-space models to combine surveys from multiple regions (or sites) into one estimate of the average long-term population growth rate and the year-to-year variability in that growth rate. Note this is not quite the same as estimating the ‘trend’; ‘trend’ often means what population change happened, whereas the long-term population growth rate refers to the underlying population dynamics. We will use as our example a dataset from harbor seals in Puget Sound, Washington, USA.


We have five regions (or sites) where harbor seals were censused from 1978-1999 while hauled out of land<sup>1</sup>. During the period of this dataset, harbor seals were recovering steadily after having been reduced to low levels by hunting prior to protection. The methodologies were consistent throughout the 20 years of the data but we do not know what fraction of the population that each region represents nor do we know the observation-error variance for each region. Given differences between behaviors of animals in different regions and the numbers of haul-outs in each region, the observation errors may be quite different. The regions have had different levels of sampling; the best sampled region has only 4 years missing while the worst has over half the years missing.

Figure 9.1 shows the data. The numbers on each line denote the different regions:

- 1 SJF
- 2 SJI
- 3 EBays
- 4 PSnd
- 5 HC

---

<sup>1</sup> Jeffries et al. 2003. Trends and status of harbor seals in Washington State: 1978-1999. *Journal of Wildlife Management* 67(1):208–219



**Fig. 9.1.** Plot of the count data from the five harbor seal regions (Jeffries et al. 2003). Each region is an index of the total harbor seal population, but the bias (the difference between the index and the true population size) for each region is unknown.

For this case study, we will assume that the underlying population process is a stochastic exponential growth process with rates of increase that were not changing through 1978-1999. However, we are not sure if all five regions sample a single “total Puget Sound” population or if there are independent subpopulations. You are going to estimate the long-term population growth rate using different assumptions about the population structures (one big population versus multiple smaller ones) and observation error structures to see how your assumptions change your estimates.

The data for this case study are in the MARSS package. The data have time running down the rows and years in the first column. We need time across the columns for the MARSS() function, so we will transpose the data:

```
dat=t(harborSealWA) #Transpose
years = dat[1,] #[1,] means row 1
n = nrow(dat)-1
dat = dat[2:nrow(dat),] #no years
```

If you needed to read data in from a comma-delimited or tab-delimited file, these are the commands to do that:

```
dat = read.csv("datafile.csv",header=TRUE)
dat = read.table("datafile.csv",header=TRUE)
```

The years (**years**) are in row 1 of **dat** and the logged data are in the rest of the rows. The number of observation time series (**n**) is the number of rows in **dat** minus 1 (for years row). Let's look at the first few years of data:

```
print(harborSealWA[1:8,], digits=3)
```

	Years	SJF	SJI	EBays	PSnd	HC
[1,]	1978	6.03	6.75	6.63	5.82	6.6
[2,]	1979	-99.00	-99.00	-99.00	-99.00	-99.0
[3,]	1980	-99.00	-99.00	-99.00	-99.00	-99.0
[4,]	1981	-99.00	-99.00	-99.00	-99.00	-99.0
[5,]	1982	-99.00	-99.00	-99.00	-99.00	-99.0
[6,]	1983	6.78	7.43	7.21	-99.00	-99.0
[7,]	1984	6.93	7.74	7.45	-99.00	-99.0
[8,]	1985	7.16	7.53	7.26	6.60	-99.0

The -99's in the data are missing values. The algorithm will ignore those values.

## 9.2 Analysis assuming a single total Puget Sound population

The first step in a state-space modeling analysis is to specify the population structure and how the regions relate to that structure. The general state-space model is

$$\mathbf{x}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{u} + \mathbf{w}_t, \text{ where } \mathbf{w}_t \sim \text{MVN}(0, \mathbf{Q}) \quad (9.1)$$

$$\mathbf{y}_t = \mathbf{Z}\mathbf{x}_t + \mathbf{a} + \mathbf{v}_t, \text{ where } \mathbf{v}_t \sim \text{MVN}(0, \mathbf{R}) \quad (9.2)$$

where all the bolded symbols are matrices. To specify the structure of the population and observations, we will specify what those matrices look like.

### 9.2.1 A single population process, $\mathbf{x}$

When we are looking at data over a large geographic region, we might make the assumption that the different census regions are measuring a single population if we think animals are moving sufficiently such that the whole area (multiple regions together) is "well-mixed". We write a model of the total population abundance as:

$$n_t = \exp(u + v_t)n_{t-1}, \quad (9.3)$$

where  $n_t$  is the total count in year  $t$ ,  $u$  is the mean population growth rate, and  $v_t$  is the deviation from that average in year  $t$ . We then take the log of both sides and write the model in log space:

$$x_t = x_{t-1} + u + v_t, \text{ where } v_t \sim N(0, q) \quad (9.4)$$

$x_t = \log n_t$ . When there is one effective population, there is one  $x$ , therefore  $\mathbf{x}_t$  is a  $1 \times 1$  matrix. There is one population growth rate ( $u$ ) and there is one process variance ( $q$ ). Thus  $\mathbf{u}$  and  $\mathbf{Q}$  are  $1 \times 1$  matrices.

### 9.2.2 The observation process, $\mathbf{y}$

For this first analysis, we assume that all five regional time series are observing this one population trajectory but they are scaled up or down relative to that trajectory. In effect, we think that animals are moving around a lot and our regional samples are some fraction of the population. There is year-to-year variation in the fraction in each region, just by chance. Notice that under this analysis, we do not think the regions represent independent subpopulations but rather independent observations of one population. Our model for the data,  $\mathbf{y}_t = \mathbf{Z}\mathbf{x}_t + \mathbf{a} + \mathbf{w}_t$ , is written as:

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} x_t + \begin{bmatrix} 0 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix} \quad (9.5)$$

Each  $y_i$  is the time series for a different region. The  $a$ 's are the bias between the regional sample and the total population. The  $a$ 's are scaling (or intercept-like) parameters<sup>2</sup>. We allow that each region could have a unique observation variance and that the observation errors are independent between regions. Lastly, we assume that the observations errors on  $\log(\text{counts})$  are normal and thus the errors on (counts) are log-normal.<sup>3</sup>

We specify independent observation errors with unique variances by specifying that the  $w$ 's come from a multivariate normal distribution with variance-covariance matrix  $\mathbf{R}$  ( $\mathbf{w} \sim \text{MVN}(0, \mathbf{R})$ ), where

<sup>2</sup> To get rid of the  $a$ 's, we scale multiple observation time series against each other; thus one  $a$  will be fixed at 0. Estimating the bias between regional indices and the total population is important for getting an estimate of the total population size. The type of time-series analysis that we are doing here (trend analysis) is not useful for estimating  $a$ 's. Instead to get  $a$ 's one would need some type of mark-recapture data. However, for trend estimation, the  $a$ 's are not important. The regional observation variance captures increased variance due to a regional estimate being a smaller sample of the total population.

<sup>3</sup> The assumption of normality is not unreasonable since these regional counts are the sum of counts across multiple haul-outs.



$$\mathbf{R} = \begin{bmatrix} r_1 & 0 & 0 & 0 & 0 \\ 0 & r_2 & 0 & 0 & 0 \\ 0 & 0 & r_3 & 0 & 0 \\ 0 & 0 & 0 & r_4 & 0 \\ 0 & 0 & 0 & 0 & r_5 \end{bmatrix} \quad (9.6)$$

$\mathbf{Z}$  specifies which observation time series,  $y_{i,1:T}$ , is associated with which population trajectory,  $x_{j,1:T}$ .  $\mathbf{Z}$  is like a look up table with 1 row for each of the  $n$  observation time series and 1 column for each of the  $m$  population trajectories. A 1 in row  $i$  column  $j$  means that observation time series  $i$  is measuring state process  $j$ . Otherwise the value in  $\mathbf{Z}_{ij} = 0$ . Since we have only 1 population trajectory, all the regions must be measuring that one population trajectory. Thus  $\mathbf{Z}$  is  $n \times 1$ :

$$\mathbf{Z} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \quad (9.7)$$

### 9.2.3 Set the constraints for MARSS

Now that we have specified our state-space model, we set the arguments that will tell the function `MARSS()` the structure of our model. We do this by passing in the argument `constraint` to `MARSS()`. `constraint` is a list which specifies any constraints on  $\mathbf{Z}$ ,  $\mathbf{u}$ ,  $\mathbf{Q}$ , etc. The function call will now look like:

```
kem = MARSS(dat, constraint=list(Z=Z.constraint, U=U.constraint,
                                  Q=Q.constraint, R=R.constraint) )
```

First we set the  $\mathbf{Z}$  constraint. We need to tell the `MARSS` function that  $\mathbf{Z}$  is a  $5 \times 1$  matrix of 1s (as in Equation 9.5). We can do this two ways. We can pass in `Z.constraint` as a matrix of ones, `matrix(1,5,1)`, just like in Equation (9.5) or we can pass in a vector of five factors, `factor(c(1,1,1,1,1))`. The  $i$ -th factor specifies which population trajectory the  $i$ -th observation time series belongs to. Since there is only one population trajectory in this first analysis, we will have a vector of five 1's: every observation time series is measuring the first, and only, population trajectory.

```
Z.constraint = factor(c(1,1,1,1,1))
```

Note, the vector (the `c()` bit) must be wrapped in `factor()` so that `MARSS` recognizes what it is. You can use either numeric or character vectors: `c(1,1,1,1,1)` is the same as `c("PS","PS","PS","PS","PS")`.


Next we specify that the  $\mathbf{R}$  variance-covariance matrix only has terms on the diagonal (the variances) with the off-diagonal terms (the covariances) equal to zero<sup>4</sup>:

<sup>4</sup> For the EM function in the `MARSS 1.0` package, the measurement errors must be uncorrelated if there are missing values in the data.

`R.constraint = "diagonal and unequal"`

The ‘and unequal’ part specifies that the variances are allowed to be unique on the diagonal. If we wanted to force the observation variances to be equal at all regions, we would use `"diagonal and equal"`.

For the first analysis, we only need to set constraints on **Z** and **R**. Since there is only one population, there is only one **u** and **Q** (they are scalars), so there are no constraints to set on them.



**Fig. 9.2.** Plot of the estimate of “log total harbor seals in Puget Sound” (minus the unknown bias for time series 1) against the data. The estimate of the total seal count has been scaled relative to the first time series. The 95% confidence intervals on the population estimates are the dashed lines. These are not the confidence intervals on the observations, and the observations (the numbers) will not fall between the confidence interval lines.

#### 9.2.4 The `MARSS()` output

The output from `MARSS()`, here assigned the name `kem`, is a list of objects. To see all the objects in it, type:

```
names(kem1)
```

The maximum-likelihood estimates of “total harbor seal population” scaled to the first observation data series (Figure 9.2) are in `kem1$states`, and `kem1$states.se` are the standard errors on those estimates. To get 95% confidence intervals, use `kem1$states +/- 1.96*kem1$states.se`. Figure 9.2 shows a plot of `kem1$states` with its 95% confidence intervals over the data. Because `kem1$states` has been scaled relative to the first time series, it is on top of that time series. One of the biases, the  $a_s$ , cannot be estimated and arbitrarily our algorithm chooses  $a_1 = 0$ , so the population estimate is scaled to the first observation time series.

The estimated parameters are a list: `kem1$par`. To get the element  $U$  of that list, which is the estimated long-term population growth rate, type in `kem1$par$U`. Multiply by 100 to get the percent increase per year. The estimated process variance is given by `kem1$par$Q`.

The log-likelihood of the fitted model is in `kem1$logLik`. We estimated one initial  $x$  ( $t = 1$ ), one process variance, one  $u$ , four  $a$ ’s, and five observation variances’s. So  $K = 12$  parameters. The AIC of this model is  $-2 \times \log\text{like} + 2K$ , which we can show by typing `kem1$AIC`.

### Example 9.1 (Fit the single population model)

*Analyze the harbor seal data using the single population model (Equations 9.4 and 9.5). The code for Example 9.1 shows you how to input data and send it to the function `MARSS()`. As you run the examples, add the estimates to the table at the end of the chapter so you can compare estimates across the examples.*

---

**Example 9.1 code**

Type `show.doc(MARSS, Case_study_2.R)` to open a file with all the example code.

```

#Read in data
dat=t(harborSealWA) #Transpose since MARSS needs time ACROSS columns
years = dat[1,]
n = nrow(dat)-1
dat = dat[2:nrow(dat),]
legendnames = (unlist(dimnames(dat)[1]))
#estimate parameters
Z.constraint = factor(c(1,1,1,1,1))
R.constraint = "diagonal and unequal"
kem1 = MARSS(dat, constraint=
  list(Z=Z.constraint, R=R.constraint))
#make figure
matplot(years, t(dat),xlab="",ylab="index of log abundance",
  pch=c("1","2","3","4","5"),ylim=c(5,9),bty="L")
lines(years,kem1$states-1.96*kem1$states.se,type="l",
  lwd=1,lty=2,col="red")
lines(years,kem1$states+1.96*kem1$states.se,type="l",
  lwd=1,lty=2,col="red")
lines(years,kem1$states,type="l",lwd=2)
title("Observations and total population estimate",cex.main=.9)
#show params
kem1$par
kem1$logLik
kem1$AIC

```

---

### 9.3 Changing the assumption about the observation variances

The variable `kem1$par$R` contains the estimates of the observation error variances. It is a matrix. Here is **R** from Example 9.1:

```

print(kem1$par$R, digits=3)

```

	SJF:1	SJI:2	EBays:3	PSnd:4	HC:5
SJF:1	0.0322	0.0000	0.0000	0.0000	0.000
SJI:2	0.0000	0.0351	0.0000	0.0000	0.000
EBays:3	0.0000	0.0000	0.0136	0.0000	0.000
PSnd:4	0.0000	0.0000	0.0000	0.0110	0.000
HC:5	0.0000	0.0000	0.0000	0.0000	0.197

Notice that the variances along the diagonal are all different—we estimated five unique observation variances. We might be able to improve the fit (relative to the number of estimated parameters) by assuming that the observation variance is equal across regions but the errors are independent. This means we estimate one observation variance instead of five. This is a fairly standard assumption for data that come from the uniform survey methodology<sup>5</sup>.

To impose this constraint, we set the **R** constraint to

```
R.constraint="diagonal and equal"
```

This tells MARSS that all the  $r$ 's along the diagonal in **R** are the same. To fit this model to the data, call MARSS() as:

```
Z.constraint = factor(c(1,1,1,1,1))
R.constraint = "diagonal and equal"
kem2 = MARSS(dat, constraint=
  list(Z=Z.constraint, R=R.constraint))
```

We estimated one initial  $x$ , one process variance, one  $u$ , four  $a$ 's, and one observation variance. So  $K = 8$  parameters. The AIC for this new model compared to the old model with five observation variances is:

```
c(kem1$AIC, kem2$AIC)

[1] -9.316477  8.849526
```

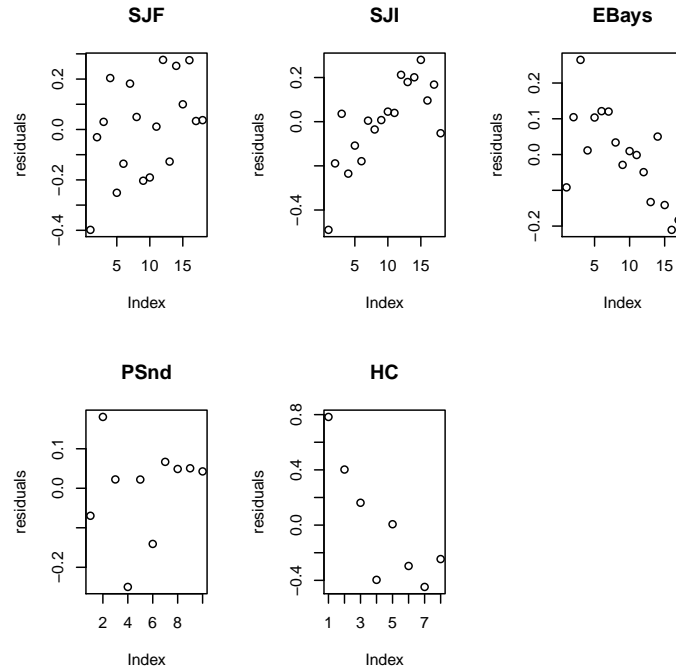
A smaller AIC means a better model. The difference between the one observation variance versus the unique observation variances is  $>10$ , suggesting that the unique observation variances model is better. Go ahead and type in the *R* code. Then add the parameter estimates to the table at the back.

One of the key diagnostics when you are comparing fits from multiple models, it to examine whether the model is flexible enough to fit the data. You do this by looking for temporal trends in the the residuals between the estimated population states (e.g. `kem2$states`) and the data. In Figure 9.3, the residuals for the second analysis are shown. Ideally, these residuals should not have a temporal trend. They should look cloud-like. The fact that the residuals have a strong temporal trend is an indication that our one population model is too restrictive for the data<sup>6</sup>.

### Example 9.2 (Fit a model with shared observation variances)

<sup>5</sup> By the way, this is not a good assumption for these data since the number haul-outs in each region varies and the regional counts are the sums across all haul-outs in a region. We will see that this is a poor assumption when we look at the AIC values.

<sup>6</sup> When comparing models via AIC, it is important that you only compare models that are flexible enough to fit the data. Fortunately if you neglect to do this, the inadequate models will usually have very high AICs and fall out of the mix anyhow.



**Fig. 9.3.** Residuals for the model with a single population. The plots of the residuals should not have trends with time, but they do... This is an indication that the single population model is inconsistent with the data. The code to make this plot is given in the script file for this case study.

*Analyze the data using the same population model as in Example 9.1, but constrain the  $\mathbf{R}$  matrix so that all five census regions have the same observation variance. The Example 9.2 code shows you how to do this. It also shows you how to make the diagnostics figure (Figure 9.3).*

**Example 9.2 code**

Type `show.doc(MARSS, Case_study_2.R)` to open a file with all the example code.

---

```
#fit model
Z.constraint = factor(c(1,1,1,1,1))
R.constraint = "diagonal and equal"
kem2 = MARSS(dat, constraint=
  list(Z=Z.constraint, R=R.constraint))
#show parameters
kem2$par$U      #population growth rate
kem2$par$Q      #process variance
kem2$par$R[1,1] #observation variance
kem2$logLik     #log likelihood
c(kem1$AIC,kem2$AIC)
#plot residuals
plotdat = t(dat); plotdat[plotdat == -99] = NA;
matrix.of.biases = matrix(kem2$par$A,
  nrow=nrow(plotdat),ncol=ncol(plotdat),byrow=T)
xs = matrix(kem2$states,
  nrow=dim(plotdat)[1],ncol=dim(plotdat)[2],byrow=F)
resids = plotdat-matrix.of.biases-xs
par(mfrow=c(2,3))
for(i in 1:n){
  plot(resids[!is.na(resids[,i]),i],ylab="residuals")
  title(legendnames[i])
}
par(mfrow=c(1,1))
```

---

## 9.4 Analysis assuming north and south subpopulations

For the third analysis, we will change our assumption about the structure of the population. We will assume that there are two subpopulations, north and south, and that regions 1 and 2 (Strait of Juan de Fuca and San Juan Islands) fall in the north subpopulation and regions 3, 4 and 5 fall in the south subpopulation. For this analysis, we will assume that these two subpopulations share their growth parameter,  $u$ , and process variance,  $q$ , since they share a similar environment and prey base. However we postulate that because of fidelity to natal rookeries for breeding, animals do not move much year-to-year between the north and south and the two subpopulations are independent.

We need to write down the state-space model to reflect this population structure. There are two subpopulations,  $x_n$  and  $x_s$ , and they have the same

growth rate  $u$ :

$$\begin{bmatrix} x_{n,t} \\ x_{s,t} \end{bmatrix} = \begin{bmatrix} x_{n,t-1} \\ x_{s,t-1} \end{bmatrix} + \begin{bmatrix} u \\ u \end{bmatrix} + \begin{bmatrix} v_{n,t} \\ v_{s,t} \end{bmatrix} \quad (9.8)$$

We specify that they are independent by specifying that their year-to-year population fluctuations (their process errors) come from a multivariate normal with no covariance:

$$\begin{bmatrix} v_{n,t} \\ v_{s,t} \end{bmatrix} \sim MVN \left( \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} q & 0 \\ 0 & q \end{bmatrix} \right) \quad (9.9)$$

For the observation process, we use the  $\mathbf{Z}$  matrix to associate the regions with their respective  $x_n$  and  $x_s$  values:

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_{n,t} \\ x_{s,t} \end{bmatrix} + \begin{bmatrix} 0 \\ a_2 \\ 0 \\ a_4 \\ a_5 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix} \quad (9.10)$$

#### 9.4.1 Specifying the MARSS() arguments

We need to change the  $\mathbf{Z}$  constraint to specify that there are two subpopulations (north and south), and that regions 1 and 2 are in the north subpopulation and regions 3,4 and 5 are in the south subpopulation. There are a few ways, we can specify this  $\mathbf{Z}$  matrix for `MARSS()`:

```
Z.constraint = matrix(c(1,1,0,0,0,0,0,1,1,1),5,2)
Z.constraint = factor(c(1,1,2,2,2))
Z.constraint = factor(c("N","N","S","S","S"))
```

Which you choose is a matter of preference as the all specify the same form for  $\mathbf{Z}$ .

We also want to specify that the  $u$ 's are the same for each subpopulation and that  $\mathbf{Q}$  is diagonal with equal  $q$ 's. To do this, we set

```
U.constraint = "equal"
Q.constraint = "diagonal and equal"
```

This says that there is one  $u$  and one  $q$  parameter and both subpopulations share it (if we wanted the  $u$ 's to be different, we would use `U.constraint="unequal"` or leave off the `u` constraint since the default behavior is `U.constraint="unequal"`).

Now we specify the new constraints and fit this model to the data:

```
Z.constraint = factor(c(1,1,2,2,2))
U.constraint = "equal"
Q.constraint = "diagonal and equal"
R.constraint = "diagonal and equal"
kem3 = MARSS(dat, constraint=list(Z=Z.constraint,
R=R.constraint, U=U.constraint, Q=Q.constraint))
```



```

Success! Parameters converged at 20 iterations.
Alert: conv.test.slope.tol is 0.5.
Test with smaller values (<0.1) to ensure convergence.

```

```

MARSS fit is
Estimation method: kem
Estimation converged in 20 iterations.
Log-likelihood: 11.52836
AIC: -7.056721   AICc: -4.73414

```

	Estimate
A.2	0.79879
A.4	-0.75590
A.5	-0.81696
Q.1	0.00812
R.1	0.02928
U.1	0.05036
x0.1	6.06203
x0.2	6.83393

```

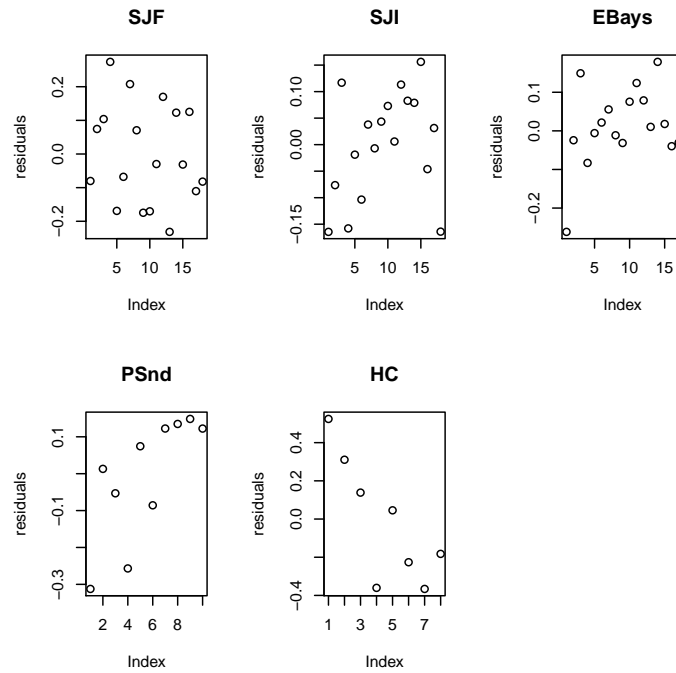
Standard errors have not been calculated.
Use MARSSparamCIs to compute CIs and bias estimates.

```

Figure 9.4 shows the residuals for the two subpopulations case. The residuals look better (more cloud-like) but the Hood Canal residuals are still temporally correlated.

### Example 9.3 (Fit a model with north and south subpopulations)

*Analyze the data using a model with two subpopulations, northern and southern. Assume that the subpopulation are independent (diagonal  $\mathbf{Q}$ ), however let each subpopulation share the same population parameters,  $u$  and  $q$ . The Example 9.3 code shows how to set the `MARSS()` arguments for this case.*



**Fig. 9.4.** The residuals for the analysis with a north and south subpopulation. The plots of the residuals should not have trends with time. Compare with the residuals for the analysis with one subpopulation.

---

### Example 9.3 code

Type `show.doc(MARSS, Case_study_2.R)` to open a file with all the example code.

```
#fit model
Z.constraint = factor(c(1,1,2,2,2))
U.constraint = "equal"
Q.constraint = "diagonal and equal"
R.constraint = "diagonal and equal"
kem3 = MARSS(dat, constraint=list(Z=Z.constraint,
  R=R.constraint, U=U.constraint, Q=Q.constraint))
#plot residuals
plotdat = t(dat); plotdat[plotdat == -99] = NA;
matrix.of.biases = matrix(kem3$par$A,
  nrow=nrow(plotdat), ncol=ncol(plotdat), byrow=T)
par(mfrow=c(2,3))
for(i in 1:n){
  j=c(1,1,2,2,2)
  xs = kem3$states[j[i],]
  resids = plotdat[,i]-matrix.of.biases[,i]-xs
  plot(resids[!is.na(resids)], ylab="residuals")
  title(legendnames[i])
}
par(mfrow=c(1,1))
```

---

## 9.5 Using MARSS() to fit other population structures

Now work through a number of different structures and fill out the table at the back of this case study. At the end you will see how your estimation of the mean population growth rate varies under different assumptions about the population and the data.

### Example 9.4 (Five subpopulations)

*Analyze the data using a model with five subpopulations, where each of the five census regions is sampling one of the subpopulations. Assume that the subpopulation are independent (diagonal  $\mathbf{Q}$ ), however let each subpopulation share the same population parameters,  $u$  and  $q$ . The Example 9.4 code shows how to set the MARSS() arguments for this case. You can use `R.constraint="diagonal and equal"` to make all the observation variances equal.*

---

#### Example 9.4 code

Type `show.doc(MARSS, Case_study_2.R)` to open a file with all the example code.

```
Z.constraint=factor(c(1,2,3,4,5))
U.constraint="equal"
Q.constraint="diagonal and equal"
R.constraint="diagonal and unequal"
kem=MARSS(dat, constraint=list(Z=Z.constraint,
    U=U.constraint, Q=Q.constraint, R=R.constraint) )
```

---

### Example 9.5 (Two subpopulations with different population parameters)

*Analyze the data using a model that assumes that the Strait of Juan de Fuca and San Juan Islands census regions represent a northern Puget Sound subpopulation, while the other three regions represent a southern Puget Sound subpopulation. This time assume that each population trajectory (north and south) has different  $u$  and  $q$  parameters:  $u_n, u_s$  and  $q_n, q_s$ . Also assume that each of the five census regions has a different observation variance. Try to write your own code. If you get stuck (or want to check your work, you can open a script file with all the Case Study 2 examples by typing `show.doc(MARSS, Case_study_2.R)` at the R command line.*

In math form, this model is:

$$\begin{bmatrix} x_{n,t} \\ x_{s,t} \end{bmatrix} = \begin{bmatrix} x_{n,t-1} \\ x_{s,t-1} \end{bmatrix} + \begin{bmatrix} u_n \\ u_s \end{bmatrix} + \begin{bmatrix} v_{n,t} \\ v_{s,t} \end{bmatrix}, \begin{bmatrix} v_{n,t} \\ v_{s,t} \end{bmatrix} \sim \text{MVN} \left( 0, \begin{bmatrix} q_n & 0 \\ 0 & q_s \end{bmatrix} \right) \quad (9.11)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_{n,t} \\ x_{s,t} \end{bmatrix} + \begin{bmatrix} 0 \\ a_2 \\ 0 \\ a_4 \\ a_5 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix} \quad (9.12)$$

**Example 9.6 (Hood Canal treated separately but covaries with others)**

Analyze the data using a model with two subpopulations with the divisions being Hood Canal versus everywhere else. In math form, this model is:

$$\begin{bmatrix} x_{p,t} \\ x_{h,t} \end{bmatrix} = \begin{bmatrix} x_{p,t-1} \\ x_{h,t-1} \end{bmatrix} + \begin{bmatrix} u_p \\ u_h \end{bmatrix} + \begin{bmatrix} v_{p,t} \\ v_{h,t} \end{bmatrix}, \begin{bmatrix} v_{p,t} \\ v_{h,t} \end{bmatrix} \sim \text{MVN} \left( 0, \begin{bmatrix} q & c \\ c & q \end{bmatrix} \right) \quad (9.13)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_{p,t} \\ x_{h,t} \end{bmatrix} + \begin{bmatrix} 0 \\ a_2 \\ a_3 \\ a_4 \\ 0 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \end{bmatrix} \quad (9.14)$$

To specify that  $\mathbf{Q}$  has one value on the diagonal (one variance) and one value on the off-diagonal (covariance) you can specify `Q.constraint` two ways:

```
Q.constraint = "equalvarcov"
Q.constraint = matrix(c("q", "c", "c", "q"), 2, 2)
```

**Example 9.7 (Three subpopulations with shared parameter values)**

Analyze the data using a model with three subpopulations as follows: north (regions 1 and 2), south (regions 3 and 4), Hood Canal (region 5). You can specify that some subpopulations share parameters while others do not. You do this by using a vector of factors for the constraints:

```

Q.constraint=factor(c("coastal","interior","interior"))
U.constraint=factor(c("puget sound","puget sound","hood canal"))
R.constraint=factor(c("boat","boat","plane","plane","plane"))

```

When `Q.constraint` and `U.constraint` are vectors (passed in as a factor), as above, they specify which  $x$ 's share parameter values. The factors must be a vector of length  $m$ , where  $m$  is the number of  $x$ 's. The  $i$ -th factor corresponds to the  $i$ -th  $x$ . In the example above, we specified that  $x_1$  has its own process variance (which we named "coastal") and  $x_2$  and  $x_3$  share a process variance value (which we named "interior"). For the long-term trends, we specified that  $x_1$  and  $x_2$  share a long-term trend ("puget sound") while  $x_3$  is allowed to have a separate trend ("hood canal").

When `R.constraint` is vector of factors, it specifies which  $y$ 's have the same observation variance. We need a  $1 \times 5$  vector here because we need to specify a value for each observation time series (there are 5). Here we imagine that observation time series 1 and 2 are boat surveys while the others are plane surveys and we want to allow the variances to differ based on methodology.

## 9.6 Discussion

There are a number of corners that we cut in order to have case study code that runs quickly:

- We ran the code starting from one initial condition. For a real analysis, you should start from a large number of random initial conditions and use the one that gives the highest likelihood. Since the EM algorithm is a "hill-climbing" algorithm, this ensures that it does not get stuck on a local maxima. `MARSS()` will do this for you if you pass it the argument `control=list(MCInit=TRUE)`. This will use a Monte Carlo routine to try many different initial conditions. See the help file on `MARSS()` for more information (by typing `?MARSS` at the *R* prompt).
- We assume independent observation and process errors. Depending on your system, observation errors may be driven by large-scale environmental factors (temperature, tides, prey locations) that would cause your observation errors to covary across regions. If your observation errors strongly covary between regions and you treat them as independent, this could be bad for your analysis. The current EM code will not handle covariance in **R** when there are missing data, but even it did, separating covariance across observation versus process errors will require much data (to have any power). In practice, the first step is to think hard about what drives sightability for your species and what are the relative levels of process and observation variance. You may be able to subsample your data in a way that will make the observation errors more independent.

- The `MARSS()` argument `control` specifies the options for the EM algorithm. We left the default tolerance for the convergence test. You would want to set this lower for a real analysis. You will need to up the `maxit` argument correspondingly.
- We used the large-sample approximation for AIC instead of a bootstrap AIC that is designed to correct for small sample size in state-space models. The bootstrap metric, AICb, takes a long time to run. Use the call `MARSSaic(kem, output=c("AICbp"))` to compute AICb. We could have shown AICc, which is the small-sample size corrector for non-state-space models. Type `kem$AICc` to get that.

Finally, in a real (maximum-likelihood) analysis, one needs to be careful not to dredge the data. The temptation is to look at the data and pick a population structure that will fit that data. This can lead to including models in your analysis that have no biological basis. In practice, we spend a lot of time discussing the population structure with biologists working on the species and review all the biological data that might tell us what are reasonable structures. From that, a set of model structures to use are selected. Other times, a particular model structure needs to be used because the population structure is not in question rather it is a matter of using that pre-specified structure and using all the data to get parameter estimates for forecasting.

## Results table

Ex.		pop. growth rate <code>par\$U</code>	process variance <code>par\$Q</code>	K <code>num. params</code>	log-like <code>logLik</code>	AIC <code>AIC</code>
1	one population different obs. vars uncorrelated					
2	one population identical obs vars uncorrelated					
3	N+S subpops identical obs vars uncorrelated;					
4	5 subpops unique obs vars <i>u</i> 's + <i>q</i> 's identical					
5	N+S subpops unique obs vars <i>u</i> 's + <i>q</i> 's identical					
6	PS + HC subpops unique obs vars <i>u</i> 's + <i>q</i> 's unique					
7	N + S + HC subpops unique obs vars <i>u</i> 's + <i>q</i> 's unique					

For AIC, only the relative differences matter. A difference of 10 between two AICs means substantially more support for the model with lower AIC. A difference of 30 or 40 between two AICs is very large.

## Questions

1. Do different assumptions about whether the observation error variances are all identical versus different affect your estimate of the long-term population growth rate (*u*)? You may want to rerun examples 3-7 with the `R.constraint` changed. `R.constraint="diagonal and unequal"` means measurement variances all different versus `"diagonal and equal"`.
2. Do assumptions about the underlying structure of the population affect your estimates of *u*? Structure here means number of subpopulations and which areas are in which subpopulation.

3. The confidence intervals for the first two analyses are very tight because the estimate process variance was very small, `kem1$par$Q`. Why do you think process variance ( $q$ ) was forced to be so small? [Hint: We are forcing there to be one and only one true population trajectory and all the observation time series have to fit that one time series. Look at the AICs too.]



## Case Study 3: Using MARSS models to identify spatial population structure and covariance


### 10.1 The problem

In this case study, we use time series of observations from nine sites along the west coast to examine large-scale spatial structure in harbor seals (Jeffries et al., 2003). Harbor seals are distributed along the west coast of the U.S. from California to Washington. The populations in Oregon and Washington have been surveyed for over 25 years at a number of haul-out sites (Figure 10.1). In general, these populations have been increasing steadily since the 1972 Marine Mammal Protection Act. It remains unknown whether they are at carrying capacity.

For management purposes, two stocks are recognized; the coastal stock consists of four sites (Northern/Southern Oregon, Coastal Estuaries, Olympic Peninsula), and the inland WA stock consists of the remaining five sites (Figure 10.1). Subtle differences exist in the demographics across sites (e.g. pupping dates), however mtDNA analyses and tagging studies have suggested that these sites may be structured on a much larger scale. Harbor seals are known for strong site fidelity, but at the same time travel large distances to forage.

Our goal for this case study is to address the following questions about spatial structure: 1) Does population abundance data support the existing management boundaries, or are there alternative groupings that receive more support? and 2) Does the Hood Canal site represent a distinct subpopulation? To address these questions, we will mathematically formulate different hypotheses about population structure via different MARSS models; each model represents a different population structure. We will then compare the data support for different models using model selection criteria, specifically AIC.

Type `show.doc(MARSS, Case_study_3.R)` to open a file with the *R* code to get you started on the analyses in this chapter.



**Fig. 10.1.** Map of spatial distribution of 9 harbor seal sites in Washington and Oregon.

## 10.2 How many distinct subpopulations?

We will analyze the support for five hypotheses about the population structure. These do not represent all possible structures but instead represent those that are considered most biologically plausible given the geography and the behavior of harbor seals.

- Hypothesis 1 Sites are grouped by stock ( $m = 2$ ), unique process variances
- Hypothesis 2 Sites are grouped by stock ( $m = 2$ ), same process variance
- Hypothesis 3 Sites are grouped by state ( $m = 2$ ), unique process variances
- Hypothesis 4 Sites are grouped by state ( $m = 2$ ), same process variance
- Hypothesis 5 All sites are part of the same panmictic population ( $m = 1$ )

Aerial survey methodology has been relatively constant across time and space, and we will assume that all sites have identical and independent observation error variance.

### 10.2.1 Specify the design, $\mathbf{Z}$ , matrices

Write down the  $\mathbf{Z}$  matrices for the hypotheses. Hint: Hypothesis 1 and 2 have the same  $\mathbf{Z}$  matrix, Hypothesis 3 and 4 have the same  $\mathbf{Z}$  matrix and Hypothesis 5 is a column of 1s.

	H 1 and 2		H 3 and 4		H 5
	$\mathbf{Z}$		$\mathbf{Z}$		$\mathbf{Z}$
	subpop	subpop	subpop	subpop	subpop
	1	2	1	2	1
Coastal Estuaries	[	]	[	]	[
Olympic Peninsula					
Str. Juan de Fuca					
San Juan Islands					
Eastern Bays					
Puget Sound					
Hood Canal					
OR North Coast					
OR South Coast					

Next you need to specify the constraints argument so that **MARSS** knows the structure of your  $\mathbf{Z}$ 's. The  $\mathbf{Z}$  constraint will be a vector of factors, i.e. it will have the form `factor(c(...))`.

- Hypothesis 1 and 2: `Z.constraint=`
- Hypothesis 3 and 4: `Z.constraint=`
- Hypothesis 5: `Z.constraint=`

### 10.2.2 Specify the grouping arguments

For this case study, we will assume that subpopulations share the same growth rate. What should `U.constraint` be for each hypothesis? To specify shared  $u$  parameters, `U.constraint` is set as a length  $m$  vector of factors and specifies which subpopulations share their  $u$  parameter. Written in *R* it takes the form `factor(c(##,...))`

- Hypothesis 1-4: `U.constraint=`
- Hypothesis 5: `U.constraint=`

What about `Q.constraint`? To specify a diagonal  $\mathbf{Q}$  matrix with shared values along the diagonal, `Q.constraint` is set as a length  $m$  vector of factors. The vector specifies which  $x_i$ 's share their process variance parameter. Look at each hypothesis (above) and write down the corresponding `Q.constraint`.

- Hypothesis 1: `Q.constraint=`
- Hypothesis 2: `Q.constraint=`
- Hypothesis 3: `Q.constraint=`
- Hypothesis 4: `Q.constraint=`
- Hypothesis 5: `Q.constraint=`

Lastly, specify `R.constraint`. As we mentioned above, we will assume that the observation errors are independent and the observation variance is the same across sites. You can specify this constraint either as a text string or as a  $n$  length vector of factors.

- Hypothesis 1-5: `R.constraint=`

### 10.2.3 Fit models and summarize results

Fit each model for each hypothesis to the seal data (look at the script `Case_Study_3.R` for the code to load the data). Each call to `MARSS()` will look like

```
kem = MARSS(sealData, constraint=list(Z = Z.constraint,
  Q = Q.constraint, R = R.constraint, U = U.constraint))
```

Fill in the following table, by fitting the five state-space models—for the five hypotheses—to the harbor seal data (using `MARSS()`). Use the `Case_Study_3.R` script so you do not have to type in all the commands.

### 10.2.4 Interpret results for question 1

What do these results indicate about the process error grouping and spatial grouping? A lower AIC means a more parsimonious model (highest likelihood given the number of parameters). A difference of 10 between AICs is large, and means the model with the higher AIC is unsupported relative to the model with lower AIC.

Extra analysis (if you have time): Do your results change if you assume that observation errors are independent but have unique variances? The nine sites have different numbers of haul-outs and so the observation variances might be different. Repeat the analysis with unique observation variances for each site (this means changing `R.constraint`). You can also try the analysis with temporally co-varying subpopulations (good and bad years correlated) by setting `Q.constraint="unconstrained"` or `Q.constraint="equalvarcov"`.

**Table 10.1.** Table to fill out for the five hypotheses from the first analysis (Section 10.2). The code of the form `$foo$bar` shows you what to type at the command line to output each parameter or metric. Remember to add the name of the model fit, e.g. `kem$foo$bar` where `kem` is the name you gave to the model fit.

H	pop. growth rate \$par\$U	proc. variance \$par\$Q	obs. variance \$par\$R	$K$ \$num. params	log-likelihood \$logLik	AIC \$AIC	AICc \$AICc
1							
2							
3							
4							
5							

### 10.3 Is Hood Canal separate?

The Hood Canal site may represent a distinct population, and has recently been subjected to a number of catastrophic events (hypoxic events, possibly leading to reduced prey availability, and several killer whale predation events, removing up to 50% of animals per occurrence). Build four models, assuming that each site (other than Hood Canal) is assigned to its current management stock, but Hood Canal is a different subpopulation ( $m = 3$ ). Again, assume that observation error variance is identical and independent across sites.

Hypothesis 1 Subpopulations have the same process variance and growth rate

Hypothesis 2 Each subpopulation has a unique process variance and growth rate

Hypothesis 3 Hood Canal has the same process variance but different growth rate

Hypothesis 4 Hood Canal has unique process variance and unique growth rate

#### 10.3.1 Specify the **Z** matrix and `Z.constraint`

The **Z** matrix for each hypothesis is the same. The coastal subpopulation consists of 4 sites (Northern/Southern Oregon, Coastal Estuaries, Olympic Peninsula), the Hood Canal subpopulation is the Hood Canal site, and the

inland WA subpopulation consists of the remaining 4 sites. Thus  $m = 3$  and  $\mathbf{Z}$  is a  $9 \times 3$  matrix:

	subpop	subpop	subpop
	1	2	3
Coastal Estuaries			
Olympic Peninsula			
Str. Juan de Fuca			
San Juan Islands			
Eastern Bays			
Puget Sound			
Hood Canal			
OR North Coast			
OR South Coast			

Then write down `Z.constraint` for this  $\mathbf{Z}$ : `factor(c(...))`

### 10.3.2 Specify which parameters are shared across which subpopulations

`U.groups` specifies which  $u$  are shared across subpopulations. Look at the hypothesis descriptions above which will specify whether subpopulations share their population growth rate or have unique population growth rates.

- Hypothesis 1: `U.constraint=`
- Hypothesis 2: `U.constraint=`
- Hypothesis 3: `U.constraint=`
- Hypothesis 4: `U.constraint=`

`U.constraint` will be a length  $m$  vector of factors. Once you have more than two subpopulations, it can get hard to keep straight which `U.constraint` goes to which subpopulation. It is best to sketch your  $\mathbf{Z}$  matrix (which tells you which site in the rows corresponds to which subpopulation in the columns). Then remember that the elements of `U.constraint` correspond one-to-one with the columns of  $\mathbf{Z}$ :

`U.constraint=factor(c(col 1 Z, col 2 Z, col 3 Z, ...))`.

Specify `Q.groups` showing which subpopulations share their process variance parameter.

- Hypothesis 1: `Q.constraint=`
- Hypothesis 2: `Q.constraint=`
- Hypothesis 3: `Q.constraint=`
- Hypothesis 4: `Q.constraint=`

`Q.constraint` will be a length  $m$  vector of factors. `R.constraint` is set so that the observation variances are the same for each site.

### 10.3.3 Fit the models and summarize results

Fit each model for each hypothesis to the seal data. Each call to `MARSS()` will look like

```
kem = MARSS(sealData, constraint=list(Z = Z.constraint,
  Q = Q.constraint, R = R.constraint, U = U.constraint))
```

**Table 10.2.** Table to fill out for the four hypotheses from the second analysis (Section 10.3). The code of the form `$foo$bar` shows you what to type at the command line to output each parameter or metric. Remember to add the name of the model fit, e.g. `kem$foo$bar` where `kem` is the name you gave to the model fit.

H	pop. growth rate \$par\$U	proc. variance \$par\$Q	obs. variance \$par\$R	$K$ \$num. params	log-likelihood \$logLik	AIC \$AIC	AICc \$AICc
1							
2							
3							
4							

### 10.3.4 Interpret results for question 2

How do the residuals for the Hood Canal site compare from these models relative to the best model from question 1? Hint: If you have the vector of estimated population states (`Xpred = t(kem$states)`) and the data (`Xobs = sealData`), the residuals for site  $i$  can be plotted in *R* as:

```
Xpred = t(kem$states)
Xobs = sealData
plot(Xpred[, Z.constraint[i]] - Xobs[,i],
  ylab="Predicted-Observed Data")
```

In *R*, if you have a matrix `Y[1:numYrs, 1:n]`, you can extract column *j* by writing `Yj = Y[,j]`.

Relative to the previous models from question 1, do these scenarios have better or worse AIC scores (smaller AIC is better)? If you were to provide advice to managers, would you recommend that the Hood Canal population is a source or sink? What implications does this have for population persistence?

---

**Code for Case Study 3**

Type `show.doc(MARSS, Case_study_3.R)` to open a file in *R* with all the example code.

---



## Case Study 4: Dynamic factor analysis (DFA) using MARSS

### 11.1 Dynamic factor analysis

In this case study, we use MARSS to do dynamic factor analysis (DFA) to look for common a set of underlying state processes that explain a set of time series (Harvey, 1989, sec. 8.5). Zuur et al. (2003) show a number of examples of using DFA to analyze fisheries catch data and zoobenthic abundances. This is conceptually different that what we have been doing in the previous case studies. Here we are trying to explain a set of  $n$  time series using linear combinations of a set of  $m$  hidden random walks.

A DFA model is a type of MARSS model with the following structure:

$$\begin{aligned} \mathbf{y}_t &= \mathbf{Z}\mathbf{x}_t + \mathbf{a} + \mathbf{v}_t \\ \text{where } \mathbf{v}_t &\sim \text{MVN}(0, \mathbf{R}) \\ \\ \mathbf{x}_t &= \mathbf{x}_{t-1} + \mathbf{w}_t \\ \text{where } \mathbf{w}_t &\sim \text{MVN}(0, \mathbf{Q}) \\ \mathbf{x}_0 &\sim \text{MVN}(\boldsymbol{\pi}, \mathbf{V}_0) \end{aligned} \tag{11.1}$$

The  $\mathbf{x}$ 's are the hidden random walks (or trends) and  $\mathbf{Z}\mathbf{x}_t$  is a linear combination of these. The objective is to estimate  $\mathbf{Z}$  and  $\mathbf{x}$ . The  $\mathbf{Z}$  matrix gives the factor loadings.

We will start by using a data set with six time series, thus  $n = 6$ . We will fit a model with three hidden random walks (or trends), thus  $m = 3$ . If we write out the model parameter matrices, this DFA model looks like so,

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \\ y_{6,t} \end{bmatrix} = \begin{bmatrix} \gamma_{11} & \gamma_{12} & \gamma_{13} \\ \gamma_{21} & \gamma_{22} & \gamma_{23} \\ \gamma_{31} & \gamma_{32} & \gamma_{33} \\ \gamma_{41} & \gamma_{42} & \gamma_{43} \\ \gamma_{51} & \gamma_{52} & \gamma_{53} \\ \gamma_{61} & \gamma_{62} & \gamma_{63} \end{bmatrix} \mathbf{x}_t + \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \\ w_{4,t} \\ w_{5,t} \\ w_{6,t} \end{bmatrix} \quad (11.2)$$

$$\begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{1,t-1} \\ x_{2,t-1} \\ x_{3,t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} v_{1,t} \\ v_{2,t} \\ v_{3,t} \end{bmatrix}$$

The  $\mathbf{B}$  and  $\mathbf{u}$  matrices have been added because MARSS will need values for these.

### 11.1.1 Constraints to ensure identifiability

If  $\mathbf{Z}$ ,  $\mathbf{A}$  and  $\mathbf{Q}$  are not constrained, then a DFA model is not identifiable (Harvey, 1989, sec 4.4). Harvey (1989, sec. 8.5.1) suggests the following constraints to make the model identifiable:

- The  $\mathbf{Q}$  matrix is constrained to be the identity matrix (a  $m \times m$  diagonal matrix of 1s).
- In the first  $m - 1$  rows of  $\mathbf{Z}$ , the  $j$ -th column in the  $i$ -th row is set to zero if  $j > i$ .
- The  $\mathbf{a}$  vector is constrained such that the first  $m$  values are zero.

Zuur et al. (2003), however, note that with Harvey's third constraint, the EM-algorithm is not particularly robust, and it takes forever to converge. Zuur et al. show that the EM estimates are much better behaved if you use a different type of constraint on the  $\mathbf{a}$  parameter. Constrain  $\mathbf{x}_t$  to have a mean of zero. This means that you replace your estimates of the hidden states,  $\mathbf{x}_t^T$  coming out of the Kalman smoother with  $\mathbf{x}_t^T - \bar{\mathbf{x}}$  for  $t = 0$  to  $T$ .  $\bar{\mathbf{x}}$  is the mean  $\mathbf{x}_t$  value across  $t$  for  $t = 0$  to  $T$  (the mean is across  $t$  not  $m$ ). With this approach, you estimate all the  $\mathbf{a}$  elements (none are set to zero) and they represent the average level of  $\mathbf{y}_t$  relative to  $\mathbf{Z}(\mathbf{x}_t - \bar{\mathbf{x}})$ . Forcing the hidden state trajectories to have a mean of zero is implemented in the `MARSS()` function using the control value `demean.states=TRUE`.

Using these constraints, the DFA model becomes

$$\mathbf{y}_t = \begin{bmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \\ y_{4,t} \\ y_{5,t} \\ y_{6,t} \end{bmatrix} = \begin{bmatrix} \gamma_{11} & 0 & 0 \\ \gamma_{21} & \gamma_{22} & 0 \\ \gamma_{31} & \gamma_{32} & \gamma_{33} \\ \gamma_{41} & \gamma_{42} & \gamma_{43} \\ \gamma_{51} & \gamma_{52} & \gamma_{53} \\ \gamma_{61} & \gamma_{62} & \gamma_{63} \end{bmatrix} \mathbf{x}_t + \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \end{bmatrix} + \begin{bmatrix} v_{1,t} \\ v_{2,t} \\ v_{3,t} \\ v_{4,t} \\ v_{5,t} \\ v_{6,t} \end{bmatrix}$$

where  $\mathbf{v}_t \sim \text{MVN}(0, \mathbf{R})$

$$\mathbf{x}_t = \begin{bmatrix} x_{1,t} \\ x_{2,t} \\ x_{3,t} \end{bmatrix} = \mathbf{I} \begin{bmatrix} x_{1,t-1} \\ x_{2,t-1} \\ x_{3,t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \\ w_{3,t} \end{bmatrix} \quad (11.3)$$

where  $\mathbf{w}_t \sim \text{MVN}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}\right)$

and  $\mathbf{x}_0 \sim \text{MVN}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 5 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 5 \end{bmatrix}\right)$

Here  $\mathbf{I}$  is the identity matrix and is a diagonal matrix with 1s on the diagonal and 0s on the off-diagonal.

## 11.2 The data

For this case study, we will analyze the Lake Washington plankton data that are loaded in the MARSS package. This is a 33-year time series of monthly counts of 13 plankton species along with temperature, phosphorous, and pH data. We want to load in the data and then extract the subset of columns that are the phytoplankton species abundances.

```
data(lakeWAplankton) #load the data
#Select out only the columns of abundance data
phytoplankton = c("Cryptomonas", "Diatom", "Green",
  "bluegreen", "Unicells", "Otheralgae")
dat.spp = lakeWAplankton[, phytoplankton]
#we will use the data from 1977 onward
dat.spp.1977 = dat.spp[181:396,]
#Transpose so time goes across columns
dat.spp.1977 = t(dat.spp.1977)
#Replace -99s with NAs to make plotting easier
dat.spp.1977[dat.spp.1977== -99] = NA
```

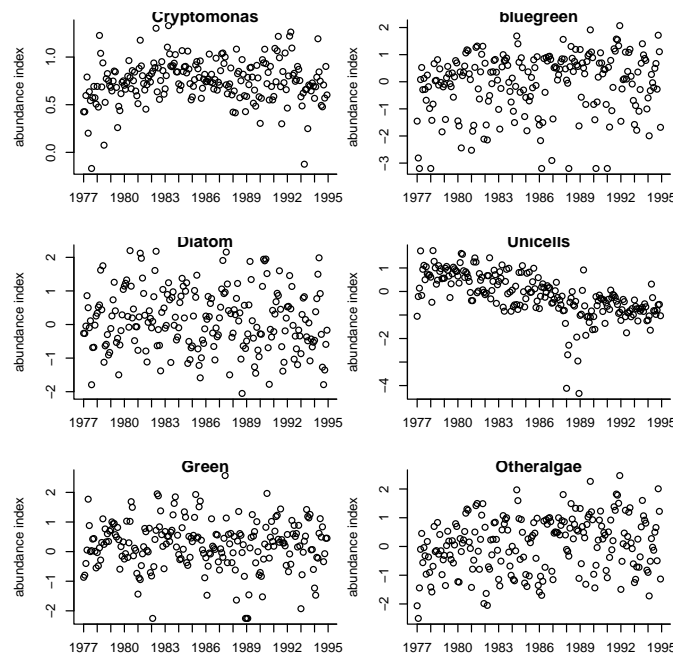
It is normal in this type of analysis to standardize each time series by subtracting its mean and dividing by its standard deviation:

$$\mathbf{y}_t^* = \Sigma^{-1}(\mathbf{y}_t - \bar{\mathbf{y}}) \quad (11.4)$$

where  $\Sigma$  is a diagonal matrix with the standard deviations of each time series along the diagonal and  $\bar{\mathbf{y}}$  is a vector of the means. In R, this can be done as follows

```
dat=dat.spp.1977
Sigma = sqrt(apply(dat,1,var,na.rm=TRUE))
Mean = apply(dat,1,mean,na.rm=TRUE)
dat.z.scores = (dat-Mean)*(1/Sigma)
rownames(dat.z.scores)=rownames(dat)
```

Figure 11.1 shows the data.



**Fig. 11.1.** Plot of the phytoplankton data (before z-scoring).

### 11.3 Setting up the model

Notice when setting up the model structure for MARSS, that the parameter matrices have a one-to-one correspondence to the model as you would write

it on paper (11.3). If a parameter matrix has a combination of fixed and estimated values, then you specify that using `matrix(list(), rows, cols)`. This is a matrix of class `list` and allows you to combine numeric and character values in a single matrix. MARSS recognizes the numeric values as fixed values and the character values as estimated values.

This is how we set up **Z** for MARSS:

```
Z.vals = list(
  "gam11", 0, 0,
  "gam21", "gam22", 0,
  "gam31", "gam32", "gam33",
  "gam41", "gam42", "gam43",
  "gam51", "gam52", "gam53",
  "gam61", "gam62", "gam63")
Z = matrix(Z.vals, 6, 3, byrow=TRUE)
```

When specifying the list values, spacing and carriage returns were added to help show the correspondence with the **Z** matrix in equation 11.3. If you print **Z** (at the R command line), you will see that it is a matrix with character values (the estimated elements) and numeric values (the fixed 0s). Notice that the 0s do not have quotes around them. If they did, it would mean the "0" is a character value and would be interpreted as the name of a parameter to be estimated not a fixed numeric value.

```
print(Z)

      [,1] [,2] [,3]
[1,] "gam11" 0    0
[2,] "gam21" "gam22" 0
[3,] "gam31" "gam32" "gam33"
[4,] "gam41" "gam42" "gam43"
[5,] "gam51" "gam52" "gam53"
[6,] "gam61" "gam62" "gam63"
```

The parameter **A** is set up similarly.

```
A.vals = list("a1", "a2", "a3", "a4", "a5", "a6")
A = matrix(A.vals, 6, 1)
```

The **Q** and **B** matrices are set to be the identity matrix using `diag()`.

```
Q = B = diag(1, 3)
```

For our first analysis, we will assume that **R** is diagonal (no covariances) and each time series of phytoplankton is allowed to have a different observation variance. Thus **R** is a diagonal matrix that looks like so:

$$\begin{bmatrix} r_{11} & 0 & 0 & 0 & 0 & 0 \\ 0 & r_{22} & 0 & 0 & 0 & 0 \\ 0 & 0 & r_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & r_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & r_{55} & 0 \\ 0 & 0 & 0 & 0 & 0 & r_{66} \end{bmatrix} \quad (11.5)$$

Each of the  $r_{ii}$  values is a different parameter to be estimated. We can specify this **R** structure using a list matrix as follows:

```
R.vals = list(
  "r11",0,0,0,0,0,
  0,"r22",0,0,0,0,
  0,0,"r33",0,0,0,
  0,0,0,"r44",0,0,
  0,0,0,0,"r55",0,
  0,0,0,0,0,"r66")
R = matrix(R.vals,6,6,byrow=TRUE)

R = "diagonal and unequal"
```

This is what **R** looks like:

```
print(R)

      [,1] [,2] [,3] [,4] [,5] [,6]
[1,] "r11" 0    0    0    0    0
[2,] 0      "r22" 0    0    0    0
[3,] 0      0      "r33" 0    0    0
[4,] 0      0      0      "r44" 0    0
[5,] 0      0      0      0      "r55" 0
[6,] 0      0      0      0      0      "r66"
```

Alternatively, MARSS has a shorthand for this structure because it is very commonly used:

```
R.vals = list(
  "r11",0,0,0,0,0,
  0,"r22",0,0,0,0,
  0,0,"r33",0,0,0,
  0,0,0,"r44",0,0,
  0,0,0,0,"r55",0,
  0,0,0,0,0,"r66")
R = matrix(R.vals,6,6,byrow=TRUE)

R = "diagonal and unequal"
```

The parameters  $\boldsymbol{\pi}$  (termed **x0** in MARSS) and **u** are set to be a matrix of zeros. Either of the following can be used:

```
x0 = U = matrix(0,3,1)
x0 = U = "zero"
```

There are shorthands for many of the common parameter matrix structures. Type ?MARSS at the R command line to see a list of the shorthand for each parameter matrix.

The  $V_0$  is a diagonal matrix with 5 along the diagonal:

```
V0 = diag(5,3)
```

Finally, we make a list of the model parameters to pass to the MARSS() function:

```
dfa.model = list( Z=Z, A=A, R=R, B=B, U=U, Q=Q, x0=x0, V0=V0)
```

## 11.4 Fitting up the model

We can now pass the DFA model list to MARSS() to estimate the  $Z$  matrix:

```
kemz.3=MARSS(dat.z.scored,constraint=dfa.model,control=list(demean.states=TRUE))
```

Success! Parameters converged at 49 iterations.

Alert: conv.test.slope.tol is 0.5.

Test with smaller values (<0.1) to ensure convergence.

MARSS fit is

Estimation method: kem

Estimation converged in 49 iterations.

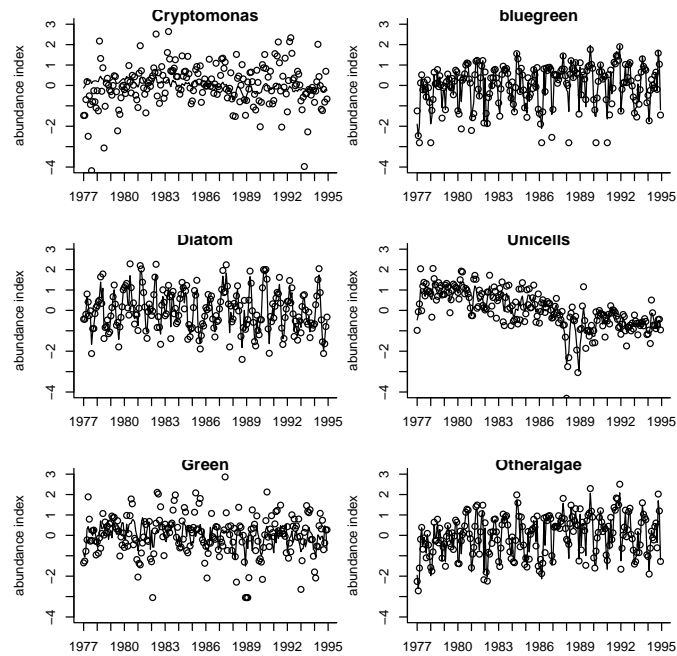
Log-likelihood: -1547.931

AIC: 3149.861 AICc: 3151.054

	Estimate
A.a1	-0.000488
A.a2	-0.001906
A.a3	-0.003471
A.a4	-0.007292
A.a5	-0.001812
A.a6	-0.007506
R.1	0.933127
R.2	0.147524
R.3	0.774248
R.4	0.121970
R.5	0.244125
R.6	0.072071
Z.gam11	0.138784
Z.gam21	0.495647
Z.gam31	0.284000

Z.gam41 0.339314  
 Z.gam51 0.437177  
 Z.gam61 0.348523  
 Z.gam22 0.616302  
 Z.gam32 -0.015477  
 Z.gam42 0.036545  
 Z.gam52 -0.120912  
 Z.gam62 0.091351  
 Z.gam33 0.356731  
 Z.gam43 0.876901  
 Z.gam53 0.044573  
 Z.gam63 0.901354

Standard errors have not been calculated.  
 Use MARSSparamCIs to compute CIs and bias estimates.



**Fig. 11.2.** Plot of the fits to the phytoplankton data.



## 11.5 Using model selection to determine the number of trends

Following (Zuur et al., 2003), we use model selection criteria (specifically AICc) to determine the number of underlying trends which have the highest data support. Our first model had three underlying trends ( $m = 3$ ). Let's compare this to a model with two underlying trends. The parameters **R** and **a** stay the same but we need to change the other parameters because  $m$  is different.

```
m=2
n=6
Q=B=diag(1,m)
x0 = U = matrix(0,m,1)
V0 = diag(5,m)
#Set up Z as a list matrix
Z = matrix(list(),n,m)
#replace all the Z values with "1" to "12" sequentially
Z[,] = as.character(1:(m*n))
#Set the correct i,j values in Z to numeric 0
for(i in 1:(m-1)) Z[i,(i+1):m]=0
dfa.model = list( Z=Z, A=A, R=R, B=B, U=U, Q=Q, x0=x0, V0=V0)

dat = dat.z.scored
kemz.2=MARSS(dat,constraint=dfa.model,control=list(demean.states=TRUE))
```

Success! Parameters converged at 30 iterations.  
Alert: conv.test.slope.tol is 0.5.  
Test with smaller values (<0.1) to ensure convergence.

MARSS fit is  
Estimation method: kem  
Estimation converged in 30 iterations.  
Log-likelihood: -1597.726  
AIC: 3241.452    AICc: 3242.32

	Estimate
A.a1	-0.00365
A.a2	-0.00234
A.a3	-0.00343
A.a4	-0.00717
A.a5	-0.00148
A.a6	-0.00758
R.1	0.64888
R.2	0.40568
R.3	0.84309

```

R.4  0.12541
R.5  0.84554
R.6  0.06974
Z.1  0.48957
Z.2  0.48553
Z.3  0.30382
Z.4  0.48356
Z.5  0.26520
Z.6  0.53216
Z.8  -0.30212
Z.9  0.27564
Z.10 0.85008
Z.11 -0.12176
Z.12 0.86334

```

Standard errors have not been calculated.  
 Use MARSSparamCIs to compute CIs and bias estimates.

```
print(c(kemz.3$AICc, kemz.2$AICc))
```

```
[1] 3151.054 3242.320
```

The **Z** matrix specification looks a little odd because we need to create a list matrix with characters and numeric values (the 0s). You need to be careful that you do not end up with a character matrix (instead list matrix) with the 0s given as "0". If you do that, MARSS will interpret the "0" as names of a **Z** parameter that needs to be estimated. You can copy and paste this code to the R command line and check that the parameters have the correct structure.

Now we can fit the model and compare the AICc values.

```

m=2
n=6
Q=B=diag(1,m)
x0 = U = matrix(0,m,1)
V0 = diag(5,m)
#Set up Z as a list matrix
Z = matrix(list(),n,m)
#replace all the Z values with "1" to "12" sequentially
Z[,] = as.character(1:(m*n))
#Set the correct i,j values in Z to numeric 0
for(i in 1:(m-1)) Z[i,(i+1):m]=0
dfa.model = list( Z=Z, A=A, R=R, B=B, U=U, Q=Q, x0=x0, V0=V0)

dat = dat.z.scored
kemz.2=MARSS(dat,constraint=dfa.model,control=list(demean.states=TRUE))

```

Success! Parameters converged at 30 iterations.  
 Alert: conv.test.slope.tol is 0.5.  
 Test with smaller values (<0.1) to ensure convergence.

MARSS fit is  
 Estimation method: kem  
 Estimation converged in 30 iterations.  
 Log-likelihood: -1597.726  
 AIC: 3241.452    AICc: 3242.32

	Estimate
A.a1	-0.00365
A.a2	-0.00234
A.a3	-0.00343
A.a4	-0.00717
A.a5	-0.00148
A.a6	-0.00758
R.1	0.64888
R.2	0.40568
R.3	0.84309
R.4	0.12541
R.5	0.84554
R.6	0.06974
Z.1	0.48957
Z.2	0.48553
Z.3	0.30382
Z.4	0.48356
Z.5	0.26520
Z.6	0.53216
Z.8	-0.30212
Z.9	0.27564
Z.10	0.85008
Z.11	-0.12176
Z.12	0.86334

Standard errors have not been calculated.  
 Use MARSSparamCIs to compute CIs and bias estimates.

```
print(c(kemz.3$AICc, kemz.2$AICc))
```

```
[1] 3151.054 3242.320
```

Now let's test a suite of models. We will test one to six underlying trends ( $m$  1 to 6) and four structures for the  $\mathbf{R}$  matrix (diagonal and unequal, diagonal and equal, unconstrained, and equal variance and equal covariance).

```

#Set up the model
levels.R = c("diagonal and unequal", "diagonal and equal",
             "unconstrained")
model.data = data.frame()
dat = dat.z.scored
n=dim(dat)[1]
for( R in levels.R ){
  for(m in 1:(n-1))
  {
    Z = matrix(list(),n,m)
    Z[,] = as.character(1:(m*n))
    if(m>1){
      for(i in 1:(m-1)) Z[i,(i+1):m]=0
    }
    x0 = U = matrix(0,m,1)
    Q = B = diag(1,m)
    V0 = diag(5,m)
    dfa.model = list( Z=Z, A=A, R=R, B=B, U=U, Q=Q, x0=x0, V0=V0)
    kemz=MARSS(dat,constraint=dfa.model,control=list(demean.states=TRUE))
    model.data=rbind(model.data,data.frame(R=R,m=m,logLik=kemz$logLik,
      AICc=kemz$AICc, K=kemz$num.params, stringsAsFactors=FALSE))
    assign(paste("kemz",m,R,sep="."),kemz)
  }
}

```

**Table 11.1.** The model fits.

	R	m	logLik	AICc	K
	unconstrained	4	-1455.45	3004.22	45.00
	unconstrained	3	-1459.99	3006.86	42.00
	unconstrained	5	-1455.24	3008.10	47.00
	unconstrained	2	-1469.74	3017.83	38.00
	diagonal and unequal	4	-1512.44	3086.36	30.00
	diagonal and unequal	5	-1532.05	3129.77	32.00
	diagonal and unequal	3	-1547.93	3151.05	27.00
	diagonal and equal	5	-1592.27	3239.73	27.00
	diagonal and unequal	2	-1597.73	3242.32	23.00
	diagonal and equal	4	-1601.50	3254.03	25.00
	diagonal and equal	3	-1614.45	3273.69	22.00
	unconstrained	1	-1611.01	3289.80	33.00
	diagonal and unequal	1	-1644.37	3325.28	18.00
	diagonal and equal	2	-1688.60	3413.73	18.00
	diagonal and equal	1	-1770.42	3567.12	13.00

## 11.6 Using a varimax rotation to determine the trend loadings

As (Harvey, 1989, p. 450) discusses in section 8.5.1, there are multiple equivalent solutions to the dynamic factor loadings. We arbitrarily constrained  $\mathbf{Z}$  in such a way to choose only one of these solutions, however the different possible solutions are equivalent. The different solutions can be related to each other by a rotation matrix  $\mathbf{H}$ . Let  $\mathbf{H}$  be any  $m \times m$  non-singular matrix. The following are equivalent solutions:

$$\begin{aligned}\mathbf{y}_t &= \mathbf{Z}\mathbf{x}_t + \mathbf{a} + \mathbf{v}_t \\ \mathbf{x}_t &= \mathbf{x}_{t-1} + \mathbf{w}_t\end{aligned}\tag{11.6}$$

and

$$\begin{aligned}\mathbf{y}_t &= \mathbf{Z}\mathbf{H}^{-1}\mathbf{x}_t + \mathbf{a} + \mathbf{v}_t \\ \mathbf{H}\mathbf{x}_t &= \mathbf{H}\mathbf{x}_{t-1} + \mathbf{H}\mathbf{w}_t\end{aligned}\tag{11.7}$$

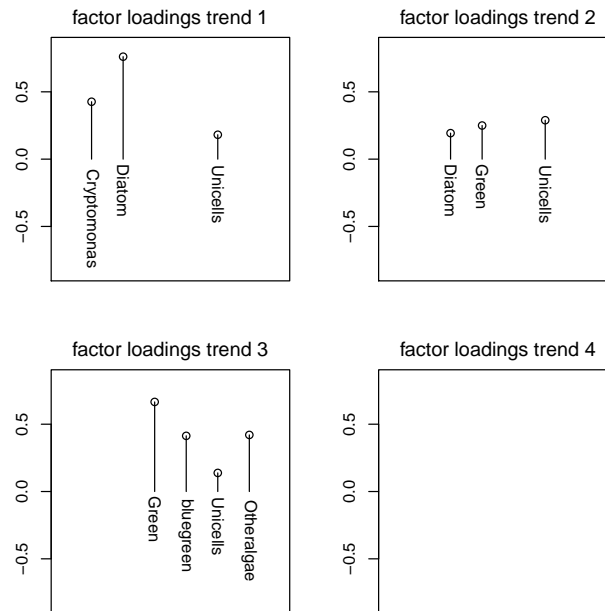
There are many ways of doing factor rotations, and a common approach is the varimax rotation which seeks a rotation matrix  $\mathbf{H}$  that tries to create the largest difference between loadings. For example, let's say there are three trends in our model. In our estimated  $\mathbf{Z}$  matrix, let's say row 3 is (0.2, 0.2, 0.2). That would mean that data series 3 is equally described by trends 1, 2, and 3. If instead row 3 was (0.8, 0.1, 0.2), this would make interpretation easier because we could say that data time series 3 was mostly described by trend 1. The varimax rotation finds the  $\mathbf{H}$  matrix that makes the  $\mathbf{Z}$  rows more like (0.8, 0.2, 0.2) and less like (0.2, 0.2, 0.2).

The varimax rotation is easy to compute in R since R has a built in function for this.

```
b=sort(model.data$AICc,index.return=TRUE)
best.model=model.data[b$ix[1],]
fitname=paste("kemz",best.model$m,best.model$R,sep=".")
best.fit=get(fitname)
Zrot = varimax(best.fit$par$Z)$loadings #Z%*%H^{-1}
```

## 11.7 Discussion

These models can take a long time to converge. Make sure to try different init values and force the algorithm to run a long time by using `control=list(minit=big, maxit=big)`, where “big” is a big number.



**Fig. 11.3.** Plot of the factor loadings (after varimax rotation) to the phytoplankton data. Factor loadings that are less than 0.1 are not plotted. Why is the trend 4 plot blank? It turns out that after the varimax rotation, the loadings on trend 4 are close to 0. The loadings on the other 3 trends are close to those from the model with  $m = 3$  (instead of  $m = 4$ ). This suggests that we choose the model with  $m = 3$  (2nd row in table) instead of  $m = 4$ .

## Case Study 5: Using state-space models to analyze noisy animal tracking data

### 12.1 A simple random walk model of animal movement

A simple random walk model of movement with drift (directional movement) but no correlation is

$$x_{1,t} = x_{1,t-1} + u_1 + v_{1,t}, \quad v_{1,t} \sim N(0, \sigma_1^2) \quad (12.1)$$

$$x_{2,t} = x_{2,t-1} + u_2 + v_{2,t}, \quad v_{2,t} \sim N(0, \sigma_2^2) \quad (12.2)$$

where  $x_{1,t}$  is the location at time  $t$  along one axis (here, longitude) and  $x_{2,t}$  is for another, generally orthogonal, axis (in here, latitude). The parameter  $u_1$  is the rate of longitudinal movement and  $u_2$  is the rate of latitudinal movement. We add errors to our observations of location:

$$y_{1,t} = x_{1,t} + w_{1,t}, \quad w_{1,t} \sim N(0, \eta_1^2) \quad (12.3)$$

$$y_{2,t} = x_{2,t} + w_{2,t}, \quad w_{2,t} \sim N(0, \eta_2^2), \quad (12.4)$$

This model is actually comprised of two separate univariate state-space models. Note that  $y_1$  depends only on  $x_1$  and  $y_2$  depends only on  $x_2$ . There are no actual interactions between these two univariate models. However, we can write the model down in the form of a multivariate model using diagonal variance-covariance matrices and a diagonal design ( $\mathbf{Z}$ ) matrix. Because the variance-covariance matrices and  $\mathbf{Z}$  are diagonal, the  $x_1:y_1$  and  $x_2:y_2$  processes will be independent as intended. Here are Equations 12.2 and 12.4 written as a MARSS model (in matrix form):

$$\begin{bmatrix} x_{1,t} \\ x_{2,t} \end{bmatrix} = \begin{bmatrix} x_{1,t-1} \\ x_{2,t-1} \end{bmatrix} + \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} + \begin{bmatrix} w_{1,t} \\ w_{2,t} \end{bmatrix}, \quad \mathbf{w}_t \sim \text{MVN} \left( 0, \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix} \right) \quad (12.5)$$

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_{1,t} \\ x_{2,t} \end{bmatrix} + \begin{bmatrix} v_{1,t} \\ v_{2,t} \end{bmatrix}, \quad \mathbf{v}_t \sim \text{MVN} \left( 0, \begin{bmatrix} \eta_1^2 & 0 \\ 0 & \eta_2^2 \end{bmatrix} \right) \quad (12.6)$$

The variance-covariance matrix for  $\mathbf{w}_t$  is a diagonal matrix with unequal variances,  $\sigma_1^2$  and  $\sigma_2^2$ . The variance-covariance matrix for  $\mathbf{v}_t$  is a diagonal matrix with unequal variances,  $\eta_1^2$  and  $\eta_2^2$ . We can write this succinctly as

$$\mathbf{x}_t = \mathbf{x}_{t-1} + \mathbf{u} + \mathbf{v}_t, \quad \mathbf{w}_t \sim \text{MVN}(0, \mathbf{Q}) \quad (12.7)$$

$$\mathbf{y}_t = \mathbf{Z}\mathbf{x}_t + \mathbf{w}_t, \quad \mathbf{v}_t \sim \text{MVN}(0, \mathbf{R}). \quad (12.8)$$

The  $\mathbf{Z}$  matrix is a  $2 \times 2$  identity matrix.

## 12.2 The problem

Loggerhead sea turtles (*Caretta caretta*) are listed as threatened under the United States Endangered Species Act of 1973. Over the last ten years, a number of state and local agencies have been deploying ARGOS tags on loggerhead turtles on the east coast of the United States. We have data on eight individuals over that period. In this case study, we use some turtle data from the WhaleNet Archive of STOP Data, however we have corrupted this data severely by adding random errors in order to create a “Bad Tag” problem. We corrupted latitude and longitude data by errors (Figure 12.1) and it would appear that our sea turtles are becoming land turtles (at least part of the time).

For this case study, we will use the `MARSS()` function to estimate true positions and speeds from the corrupted data. We will use a mapping package to plot the results: the `maps` package. If you have not already, install this package by selecting the ‘Packages’ menu and then ‘Install packages’ and then select `maps`. If you are on a Mac, remember to select “binaries” for the package type. Type `show.doc(MARSS, Case_study_5.R)` to open a file in *R* with all *R* code to get you started on the analyses in this chapter.


## 12.3 Estimate locations from bad tag data

### 12.3.1 Read in the data and load maps package

Our noisy data are in `loggerheadNoisy`. They consist of daily readings of location (longitude/latitude). The data are recorded daily and `MARSS()` requires a data entry for each day. If data are missing for a day, then the entries for latitude and longitude for that day should be -99. However, to make this case study run quickly, we have interpolated all missing values in the original, uncorrupted, dataset (`loggerhead`). The corrupted data are a dataframe and the first six lines look like so

```
loggerheadNoisy[1:6,]
```





**Fig. 12.1.** Plot of the tag data from the turtle Big Mama. Errors in the location data make it seem that Big Mama has been moving overland.

	turtle	month	day	year	lon	lat
1	BigMama	5	28	2001	-81.45989	31.70337
2	BigMama	5	29	2001	-80.88292	32.18865
3	BigMama	5	30	2001	-81.27393	31.67568
4	BigMama	5	31	2001	-81.59317	31.83092
5	BigMama	6	1	2001	-81.35969	32.12685
6	BigMama	6	2	2001	-81.15644	31.89568

The file has data for eight turtles:

```
levels(loggerheadNoisy$turtle)

[1] "BigMama" "Bruiser" "Humpty" "Isabelle" "Johanna"
[6] "MaryLee" "TBA" "Yoto"
```

We will first analyze the position data for “Big Mama”. We put the data for “Big Mama” into variable `dat`. `dat` is transposed because we need time across the columns.

```

turtlename="BigMama"
dat = loggerheadNoisy[which(loggerheadNoisy$turtle==turtlename),5:6]
dat = t(dat) #transpose

```

We will begin by specifying the structure of the MARSS model and then use `MARSS()` to fit that model to the data. There are two state processes (one for latitude and the other for longitude), and there is one observation time series for each state process. As we saw in Equation 12.6,  $\mathbf{Z}$  is the an identity matrix (a diagonal matrix with 1s on the diagonal). We could specify this structure as `Z.constraint="identity"` or `Z.constraint=factor(c(1,2))`. Although technically, this is unnecessary as this is the default constraint for  $\mathbf{Z}$ .

We will assume that the errors are independent and that there are different drift rates ( $u$ ), process variances ( $\sigma^2$ ) and observation variances for latitude and longitude ( $\eta^2$ ).

```

Z.constraint=factor(c(1,2))
U.constraint="unequal"
Q.constraint="diagonal and unequal"
R.constraint="diagonal and unequal"

```

Fit the model to the data:

```

kem = MARSS(dat, constraint=list(Z = Z.constraint,
                                Q = Q.constraint, R = R.constraint, U = U.constraint))

```

### 12.3.2 Compare state estimates to the real positions

The real locations (from which `loggerheadNoisy` was produced by adding noise) are in `loggerhead`. In Figure 12.2, we compare the tracks estimated from the noisy data with the original, good, data (see `Case_Study_5.R` for the code to make this plot. There are only a few data points for the real data because in the real tag data, there are many missing days.


### 12.3.3 Estimate speeds for each turtle

Turtle biologists designated one of these loggerheads “Big Mama,” presumably for her size and speed. For each of the eight turtles, estimate the average miles traveled per day. To calculate the distance traveled by a turtle each day, you use the estimate (from `MARSS()`) of the lat/lon location of turtle at day  $t$  and at day  $t - 1$ . To calculate distance traveled in miles from lat/lon start and finish locations, we will use the function `GCDF` defined at the beginning of the `R` script, `Case_Study_5.R`):

```

distance[i-1]=GCDF(pred.lon[i-1],pred.lon[i],
                   pred.lat[i-1],pred.lat[i])

```




**Fig. 12.2.** Plot of the estimated track of the turtle Big Mama versus the good location data (before we corrupted it with noise).

`pred.lon` and `pred.lat` are the predicted longitudes and latitudes from `MARSS()`: rows one and two in `kem$states`. To calculate the distances for all days, we put this through a `for` loop:

```
distance = array(-99, dim=c(dim(dat)[2]-1,1))
for(i in 2:dim(dat)[2])
  distance[i-1]=GCDF(pred.lon[i-1],pred.lon[i],
    pred.lat[i-1],pred.lat[i])
```

The command `mean(distance)` gives us the average distance per day. We can also make a histogram of the distances traveled per day (Figure 12.3). Repeat the analysis done for “Big Mama” for each of the other turtles and fill

out the speed table (Table 12.1).



**Fig. 12.3.** Histogram of the miles traveled per day for Big Mama. Compare this to the estimate of miles traveled per day if you had not accounted for measurement errors. See the script file, `Case_Study_5.R`, for the code to do this.

## 12.4 Comparing turtle tracks to proposed fishing areas

One of the greatest threats to the long term viability of loggerhead turtles is incidental take by net/pot fisheries. Add two proposed fishing areas to your turtle plots:

```
# the proposed fishery areas
lines(c(-77,-78,-78,-77,-77),
      c(33.5,33.5,32.5,32.5,33.5),col="red",lwd=2)
lines(c(-75,-76,-76,-75,-75),
      c(38,38,37,37,38),col="red",lwd=2)
```

Given that only one area can be chosen as a future fishery, what do your predicted movement trajectories for our eight turtles tell you?

**Table 12.1.** Estimated speeds with location errors included in model versus speeds when we assume that the data have no location error.

Turtle	Location error included in model	Data assumed to be error free
Big Mama		
Bruiser		
Humpty		
Isabelle		
Johanna		
Mary Lee		
TBA		
Yoto		

## 12.5 Using specialized packages to analyze tag data

If you have real tag data to analyze, you should use a state-space modeling package that is customized for fitting MARSS models to tracking data. The MARSS package does not have all the bells and whistles that you would want for analyzing tracking data, particularly tracking data in the marine environment. These are a couple *R* packages that we have come across for this purpose:

UKFSST <http://www.soest.hawaii.edu/tag-data/tracking/ukfsst/>

KFTRACK <http://www.soest.hawaii.edu/tag-data/tracking/kftrack/>

**kftrack** is a full-featured toolbox for analyzing tag data with extended Kalman filtering. It incorporates a number of extensions that are important for analyzing track data: barriers to movement such as coastlines and non-Gaussian movement distributions. With **kftrack**, you can use the real tag data which has big gaps, i.e. days with no location. **MARSS()** will struggle with these data because it will estimate states for all the unseen days; **kftrack** only fits to the seen days.

To use **kftrack** to fit the turtle data, type

```
library(kftrack) # must be installed from a local zip file
loggerhead = loggerhead
# Run kftrack with the first turtle (BigMama)
```

```
turtlename = "BigMama"  
dat = loggerhead[which(loggerhead$turtle == turtlename),2:6]  
model = kftrack(dat, fix.first=F, fix.last=F,  
                var.struct="uniform")
```

## A

---

### **Textbooks and articles that use MARSS modeling for population modeling**

#### **Textbooks Describing the Estimation of Process and Non-process Variance**

There are many textbooks on Kalman filtering and estimation of state-space models. The following are a sample of books on state-space modeling that we have found especially helpful.

Shumway, R. H., and D. S. Stoffer. 2006. Time series analysis and its applications. Springer-Verlag, New York, New York, USA.

Harvey, A. C. 1989. Forecasting, structural time series models and the Kalman filter. Cambridge University Press, Cambridge, UK.

Durbin, J., and S. J. Koopman. 2001. Time series analysis by state space methods. Oxford University Press, Oxford.

Kim, C. J. and Nelson, C. R. (1999), State space models with regime switching, MIT Press, Cambridge, Massachusetts.

King, R., G. Olivier, B. Morgan, and S. Brooks. 2009. Bayesian analysis for population ecology.

Giovanni, P., S. Petrone, and P. Campagnoli. 2009. Dynamic linear models in R.

Pole, A., M. West, and J. Harrison. 1994. Applied Bayesian forecasting and time series analysis.

Bolker, B. 2008. Ecological models and data in R. Princeton University Press.

West, M. and Harrison, J. (1997), Bayesian forecasting and dynamic models, Springer-Verlag.

#### **Maximum-likelihood papers**

This is just a sample of the papers from the population modeling literature.

de Valpine, P. 2002. Review of methods for fitting time-series models with process and observation error and likelihood calculations for nonlinear, non-Gaussian state-space models. *Bulletin of Marine Science* 70:455-471.

de Valpine, P. and A. Hastings. 2002. Fitting population models incorporating process noise and observation error. *Ecological Monographs* 72:57-76.

de Valpine, P. 2003. Better inferences from population-dynamics experiments using Monte Carlo state-space likelihood methods. *Ecology* 84:3064-3077.

de Valpine, P. and R. Hilborn. 2005. State-space likelihoods for nonlinear fisheries time series. *Canadian Journal of Fisheries and Aquatic Sciences* 62:1937-1952.

Dennis, B., J.M. Ponciano, S.R. Lele, M.L. Taper, and D.F. Staples. 2006. Estimating density dependence, process noise, and observation error. *Ecological Monographs* 76:323-341.

Ellner, S.P. and E.E. Holmes. 2008. Resolving the debate on when extinction risk is predictable. *Ecology Letters* 11:E1-E5.

Erzini, K. 2005. Trends in NE Atlantic landings (southern Portugal): identifying the relative importance of fisheries and environmental variables. *Fisheries Oceanography* 14:195-209.

Erzini, K., Inejih, C. A. O., and K. A. Stobberup. 2005. An application of two techniques for the analysis of short, multivariate non-stationary time-series of Mauritanian trawl survey data ICES Journal of Marine Science 62:353-359.

Hinrichsen, R.A. and E.E. Holmes. 2009. Using multivariate state-space models to study spatial structure and dynamics. In *Spatial Ecology* (editors Robert Stephen Cantrell, Chris Cosner, Shigui Ruan). CRC/Chapman Hall.

Hinrichsen, R.A. 2009. Population viability analysis for several populations using multivariate state-space models. *Ecological Modelling* 220:1197-1202.

Holmes, E.E. 2001. Estimating risks in declining populations with poor data. *Proceedings of the National Academy of Sciences of the United States of America* 98:5072-5077.

Holmes, E.E. and W.F. Fagan. 2002. Validating population viability analysis for corrupted data sets. *Ecology* 83:2379-2386.

Holmes, E.E. 2004. Beyond theory to application and evaluation: diffusion approximations for population viability analysis. *Ecological Applications* 14:1272-1293.

Holmes, E.E., W.F. Fagan, J.J. Rango, A. Folarin, S.J.A., J.E. Lippe, and N.E. McIntyre. 2005. Cross validation of quasi-extinction risks from real time series: An examination of diffusion approximation methods. U.S. Department of Commerce, NOAA Tech. Memo. NMFS-NWFSC-67, Washington, DC.

Holmes, E.E., J.L. Sabo, S.V. Viscido, and W.F. Fagan. 2007. A statistical approach to quasi-extinction forecasting. *Ecology Letters* 10:1182-1198.

Kalman, R.E. 1960. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82:35-45.



Lele, S.R. 2006. Sampling variability and estimates of density dependence: a composite likelihood approach. *Ecology* 87:189-202.

Lele, S.R., B. Dennis, and F. Lutscher. 2007. Data cloning: easy maximum likelihood estimation for complex ecological models using Bayesian Markov chain Monte Carlo methods. *Ecology Letters* 10:551-563.

Lindley, S.T. 2003. Estimation of population growth and extinction parameters from noisy data. *Ecological Applications* 13:806-813.

Ponciano, J.M., M.L. Taper, B. Dennis, S.R. Lele. 2009. Hierarchical models in ecology: confidence intervals, hypothesis testing, and model selection using data cloning. *Ecology* 90:356-362.

Staples, D.F., M.L. Taper, and B. Dennis. 2004. Estimating population trend and process variation for PVA in the presence of sampling error. *Ecology* 85:923-929.

Zuur, A. F., and G. J. Pierce. 2004. Common trends in Northeast Atlantic squid time series. *Journal of Sea Research* 52:57-72.

Zuur, A. F., I. D. Tuck, and N. Bailey. 2003. Dynamic factor analysis to estimate common trends in fisheries time series. *Canadian Journal of Fisheries and Aquatic Sciences* 60:542-552.

Zuur, A. F., R. J. Fryer, I. T. Jolliffe, R. Dekker, and J. J. Beukema. 2003. Estimating common trends in multivariate time series using dynamic factor analysis. *Environmetrics* 14:665-685.

## Bayesian papers

This is a sample of the papers from the population modeling and animal tracking literature.

Buckland, S.T., K.B. Newman, L. Thomas and N.B. Koestersa. 2004. State-space models for the dynamics of wild animal populations. *Ecological modeling* 171:157-175.

Calder, C., M. Lavine, P. Müller, J.S. Clark. 2003. Incorporating multiple sources of stochasticity into dynamic population models. *Ecology* 84:1395-1402.

Chaloupka, M. and G. Balazs. 2007. Using Bayesian state-space modelling to assess the recovery and harvest potential of the Hawaiian green sea turtle stock. *Ecological Modelling* 205:93-109.

Clark, J.S. and O.N. Bjørnstad. 2004. Population time series: process variability, observation errors, missing values, lags, and hidden states. *Ecology* 85:3140-3150.

Jonsen, I.D., R.A. Myers, and J.M. Flemming. 2003. Meta-analysis of animal movement using state space models. *Ecology* 84:3055-3063.

Jonsen, I.D., J.M. Flemming, and R.A. Myers. 2005. Robust state-space modeling of animal movement data. *Ecology* 86:2874-2880.

Meyer, R. and R.B. Millar. 1999. BUGS in Bayesian stock assessments. *Can. J. Fish. Aquat. Sci.* 56:1078-1087.

Meyer, R. and R.B. Millar. 1999. Bayesian stock assessment using a state-space implementation of the delay difference model. *Can. J. Fish. Aquat. Sci.* 56:37-52.

Meyer, R. and R.B. Millar. 2000. Bayesian state-space modeling of age-structured data: fitting a model is just the beginning. *Can. J. Fish. Aquat. Sci.* 57:43-50.

Newman, K.B., S.T. Buckland, S.T. Lindley, L. Thomas, and C. Fernández. 2006. Hidden process models for animal population dynamics. *Ecological Applications* 16:74-86.

Newman, K.B., C. Fernández, L. Thomas, and S.T. Buckland. 2009. Monte Carlo inference for state-space models of wild animal populations. *Biometrics* 65:572-583

Rivot, E., E. Prévost, E. Parent, and J.L. Baglinière. 2004. A Bayesian state-space modelling framework for fitting a salmon stage-structured population dynamic model to multiple time series of field data. *Ecological Modeling* 179:463-485.

Schnute, J.T. 1994. A general framework for developing sequential fisheries models. *Canadian J. Fisheries and Aquatic Sciences* 51:1676-1688.

Swain, D.P., I.D. Jonsen, J.E. Simon, and R.A. Myers. 2009. Assessing threats to species at risk using stage-structured state-space models: mortality trends in skate populations. *Ecological Applications* 19:1347-1364.

Thogmartin, W.E., J.R. Sauer, and M.G. Knutson. 2004. A hierarchical spatial model of avian abundance with application to cerulean warblers. *Ecological Applications* 14:1766-1779.

Trenkel, V.M., D.A. Elston, and S.T. Buckland. 2000. Fitting population dynamics models to count and cull data using sequential importance sampling. *J. Am. Stat. Assoc.* 95:363-374.

Viljugrein, H., N.C. Stenseth, G.W. Smith, and G.H. Steinbakk. 2005. Density dependence in North American ducks. *Ecology* 86:245-254.

Ward, E.J., R. Hilborn, R.G. Towell, and L. Gerber. 2007. A state-space mixture approach for estimating catastrophic events in time series data. *Can. J. Fish. Aquat. Sci., Can. J. Fish. Aquat. Sci.* 64:899-910.

Wikle, C.K., L.M. Berliner, and N. Cressie. 1998. Hierarchical Bayesian space-time models. *Journal of Environmental and Ecological Statistics* 5:117-154

Wikle, C.K. 2003. Hierarchical Bayesian models for predicting the spread of ecological processes. *Ecology* 84:1382-1394.

## B

---

### Package MARSS: Object structures

#### B.1 Model objects: class `marssm`

Objects of class ‘`marssm`’ specify Multivariate Autoregressive State Space (MARSS) models. The `model` component of an ML estimation object (class `marssMLE`; see below) is a `marssm` object. These objects have the following components:

**data** An optional matrix (not dataframe), in which each row is a time series (time across columns).

**fixed** A list with 8 matrices `Z`, `A`, `R`, `B`, `U`, `Q`, `x0`, `V0`, specifying which elements of each parameter are fixed.

**free** A list with 8 matrices `Z`, `A`, `R`, `B`, `U`, `Q`, `x0`, `V0`, specifying which elements of each parameter are to be estimated.

**M** An array of dim  $n \times n \times T$  (an  $n \times n$  missing values matrix for each time point). Each matrix is diagonal with 0 at the  $i, i$  value if the  $i$ -th value of `y` is missing, and 1 otherwise.

**miss.value** Specifies missing value representation in the data.

The matrices in **fixed** and **free** work as pairs to specify the fixed and free elements for each parameter. See Chapter 4. The dimensions for **fixed** and **free** matrices are as follows, where  $n$  is the number of observation time series and  $m$  is the number of state processes:

**Z**  $n \times m$

**B**  $m \times m$

**U**  $m \times 1$

**Q**  $m \times m$

**A**  $n \times 1$

**R**  $n \times n$

**x0**  $m \times 1$

**V0**  $m \times m$

Use `is.marssm()` to check whether an `marssm` object is correctly specified. The MARSS package includes an `as.marssm()` method to convert objects of class `popWrap` (see next section) to objects of class `marssm`.

## B.2 Wrapper objects: class `popWrap`

Wrapper objects of class `popWrap` contain specifications and options for estimation of a MARSS model. A `popWrap` object has the following components:

- data** A matrix (not dataframe) of observations (rows)  $\times$  time (columns).
- m** Number of hidden state processes (number of rows in  $\mathbf{x}$ ).
- constraint** Either a list with 8 string elements Z, A, R, B, U, Q, x0, V0 (see below for details), or string "use fixed/free".
- fixed** If `constraint[[elem]]="use fixed/free"`, a list with 8 matrices Z, A, R, B, U, Q, x0, V0.
- free** If `constraint[[elem]]="use fixed/free"`, a list with 8 matrices Z, A, R, B, U, Q, x0, V0.
- inits** A list with 8 matrices Z, A, R, B, U, Q, x0, V0, specifying initial values for parameters. Dimensions are given in the class 'marssm' section.
- miss.value** Specifies missing value representation (default is -99).
- method** The method used for estimation: 'kem' for Kalman-EM, 'BFGS' for quasi-Newton.
- control** List of estimation options. For the EM algorithm, these include the elements `minit`, `maxit`, `abstol`, `allow.degen`, `safe` and `trace`. For Monte Carlo initialization, these include the elements `MCInit`, `numInits`, `numInitSteps` and `boundsInits`. See class `marssMLE` section for details.

Component **constraint** is a convenient way to specify model structure for certain common cases. If `constraint="use fixed/free"`, both **fixed** and **free** must be provided. See the class `marssm` section for how to specify fixed and free matrices. The function `MARSS()` calls `popWrap()` to create a `popWrap` object, then `is.marssm()` to coerce this object to class `marssm` for the estimation function.

The `popWrap()` function calls `checkPopWrap()` to check user inputs from `MARSS()`. Valid constraints are below.

- A** May be either the string 'scaling' or the string 'zero' to specify a column vector of zeros ( $\mathbf{a} = 0$ ).
- B** String 'identity' or a numeric matrix specifying a fixed **B** matrix. The string 'zero' may be used to specify a  $m \times m$  matrix of zeros ( $\mathbf{B} = 0$ ).
- Q** String 'unconstrained', 'diagonal and unequal', 'diagonal and equal', or 'equalvarcov'. May also be numeric or character vector of class factor specifying shared diagonal values or a numeric matrix specifying a fixed **Q** matrix.

- R** String ‘unconstrained’, ‘diagonal and unequal’, ‘diagonal and equal’, or ‘equalvarcov’. May also be numeric or character vector of class factor specifying shared diagonal values or a numeric matrix specifying a fixed **R** matrix.
- U** String ‘unconstrained’=‘unequal’, or ‘equal’. May also be numeric or character vector of class factor specifying shared **u** elements or a  $m \times 1$  numeric matrix specifying a fixed **u** matrix. The string ‘zero’ may be used to specify a column vector of zeros (**u** = 0).
- x0** String ‘unconstrained’=‘unequal’, or ‘equal’. May also be vector of class factor specifying shared  $\pi$  ( $t = 1$ ) values or a  $m \times 1$  numeric matrix specifying a fixed  $\pi$  ( $t = 1$ ) matrix. The string ‘zero’ may be used to specify a column vector of zeros ( $\pi = 0$ ).
- Z** A vector of class factor specifying which **y** time series correspond to which state time series (in **x**) or a numeric  $n \times m$  matrix specifying the **Z** matrix. The string ‘identity’ can be used to specify a  $n \times n$  identity matrix and string ‘ones’ may be used to specify a column vector of  $n$  ones.

### B.3 ML estimation objects: class marssMLE

Objects of class marssMLE specify maximum-likelihood estimation for a MARSS model, both before and after fitting. A minimal marssMLE object contains components **model**, **start** and **control**, which must be present for estimation by functions like **MARSSkem()**.

- model** MARSS model specification (an object of class ‘marssm’).
- start** List with 7 matrices **A**, **R**, **B**, **U**, **Q**, **x0**, **V0**, specifying initial values for parameters. Dimensions are given in the class marssm section.
- control** A list specifying estimation options. For **method="kem"**, these are
  - minit* Minimum number of iterations in the maximization algorithm.
  - maxit* Maximum number of iterations in the maximization algorithm.
  - abstol* Optional tolerance for log-likelihood change. If log-likelihood decreases less than this amount relative to the previous iteration, the EM algorithm exits.
  - trace* A positive integer. If not zero, a record will be created of each variable the maximization iterations. The information recorded depends on the maximization method.
  - safe* If TRUE, **MARSSkem()** will rerun **MARSSkf()** after each individual parameter update rather than only after all parameters are updated.
  - MCInit* Use Monte Carlo initialization?
  - numInits* Number of random initial value draws.
  - numInitSteps* Number of iterations for each initial value draw.
  - boundsInits* Bounds on the uniform distributions from which initial values will be drawn. (Note that bounds for the covariance matrices **Q** and **R**, which require positive values, are specified in logs.)

*silent* Suppresses printing of progress bar and convergence information.

`MARSSkem()` appends the following components to the `marssMLE` object:

**method** A string specifying the estimation method ('kem' for estimation by `MARSSkem()`).

**par** A list with 8 matrices of estimated parameter values Z, A, R, B, U, Q, x0, V0. If there are fixed elements in the matrices, the corresponding elements in `$par` are set to the fixed values.

**kf** A list containing Kalman filter/smoothing output. See Chapter 2

**numIter** Number of iterations required for convergence.

**convergence** Convergence status.

**logLik** the exact Log-likelihood. See Section 5.2.

**errors** any error messages

**iter.record** record of the parameter values at each iteration (if `control$trace=1`)

Several functions append additional components to the 'marssMLE' object passed in. These include:

`MARSSaic` Appends `AIC`, `AICc`, `AICbb`, `AICbp`, depending on the AIC flavors requested.

`MARSShessian` Appends `Hessian`, `gradient`, `parMean` and `parSigma`.

`MARSSparamCIs` Appends `par.se`, `par.bias`, `par.upCI` and `par.lowCI`.

## C

---

### Package MARSS: The top-level MARSS functions and the base functions

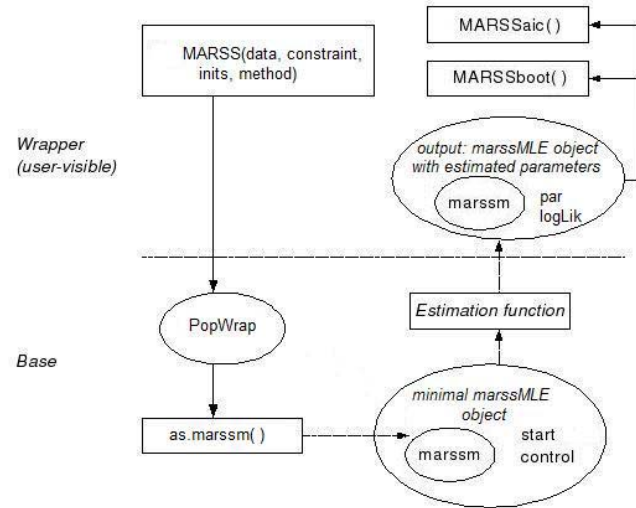
Package MARSS includes functions for estimating Multivariate Autoregressive State Space models, obtaining confidence intervals for parameters, and calculating Akaike's Information Criterion (AIC) for model selection. To make the package both flexible and easy to use, it is designed in two levels. At the base level, the programmer can interact directly with the estimation functions, using two kinds of R objects: objects of the model specification class 'marssm', and objects of estimation classes such as 'marssMLE'. At the user level, the **MARSS()** function allows model estimation with just one function call, hiding the details for ease of use. Users create models in an intuitive way by specifying constraints; the **MARSS()** function then converts these constraints into the object structures required by the estimation functions, performing error checking as necessary.

The two-level package structure allows new users convenient access to the underlying functions, while maintaining flexibility to incorporate different applications and algorithms. Developers can use the base object types to write new functions for their own modeling applications.

To use **MARSS()**, the user specifies a model by supplying the constraint argument to **MARSS()**, using the method argument to specify an estimation method. Optionally, the user may provide initial values for the free parameters, and specify estimation options; for details see the **MARSS()** help file. The function returns an object containing the model, parameter values and estimation details. The user may pass the returned object to **MARSSboot()**, which generates bootstrap parameter estimates, or to **MARSSaic()**, which calculates various versions of AIC for model selection.

Figure C.1 shows the underlying base level operations **MARSS()** performs. The function creates a wrapper object of class 'popWrap'. It then calls the **as.marssm()** method for 'popWrap' to create a **marssm** object from the constraints provided. This model object, initial values and control information are the minimal information required by the estimation functions, and are combined into an object of class appropriate for the estimation method. The estimation function adds to this object the estimated parameter values, esti-

mation details, and other function-specific components, and then returns the augmented object.



**Fig. C.1.** Two-level structure of the MARSS package. Rectangles represent functions; ovals represent objects.



---

## References

- BIERNACKI, C., CELEUX, G., AND GOVAERT, G. 2003. Choosing starting values for the EM algorithm for getting the highest likelihood in multivariate gaussian mixture models. *Computational Statistics and Data Analysis* 41:561–575.
- BROCKWELL, P. J. AND DAVIS, R. A. 1991. Time series: theory and methods. Springer-Verlag, New York, NY.
- CAVANAUGH, J. AND SHUMWAY, R. 1997. A bootstrap variant of AIC for state-space model selection. *Statistica Sinica* 7:473–496.
- DEMPSTER, A., LAIRD, N., AND RUBIN, D. 1977. Likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B* 39:1–38.
- DENNIS, B., MUNHOLLAND, P. L., AND SCOTT, J. M. 1991. Estimation of growth and extinction parameters for endangered species. *Ecological Monographs* 61:115–143.
- DENNIS, B., PONCIANO, J. M., LELE, S. R., TAPER, M. L., AND STAPLES, D. F. 2006. Estimating density dependence, process noise, and observation error. *Ecological Monographs* 76:323–341.
- ELLNER, S. P. AND HOLMES, E. E. 2008. Resolving the debate on when extinction risk is predictable. *Ecology Letters* 11:E1–E5.
- GERBER, L. R., MASTER, D. P. D., AND KAREIVA, P. M. 1999. Grey whales and the value of monitoring data in implementing the u.s. endangered species act. *Conservation Biology* 13:1215–1219.
- GHAHRAMANI, Z. AND HINTON, G. E. 1996. Parameter estimation for linear dynamical systems. Technical Report CRG-TR-96-2, University of Toronto, Dept. of Computer Science.
- HARVEY, A. C. 1989. Forecasting, structural time series models and the Kalman filter. Cambridge University Press, Cambridge, UK.
- HARVEY, A. C. AND SHEPHARD, N. 1993. Structural time series models. In G. S. Maddala, C. R. Rao, and H. D. Vinod (eds.), *Handbook of Statistics*, Volume 11. Elsevier Science Publishers B V, Amsterdam.

- HINRICHSEN, R. 2009. Population viability analysis for several populations using multivariate state-space models. *Ecological Modelling* 220:1197–1202.
- HINRICHSEN, R. AND HOLMES, E. E. 2009. Using multivariate state-space models to study spatial structure and dynamics. In R. S. Cantrell, C. Cosner, and S. Ruan (eds.), *Spatial Ecology*. CRC/Chapman Hall.
- HOLMES, E. E. 2001. Estimating risks in declining populations with poor data. *Proceedings of the National Academy of Sciences of the United States of America* 98:5072–5077.
- HOLMES, E. E. 2004. Beyond theory to application and evaluation: diffusion approximations for population viability analysis. *Ecological Applications* 14:1272–1293.
- HOLMES, E. E. 2010. Derivation of the EM algorithm for constrained and unconstrained marss models. Technical report, Northwest Fisheries Science Center, Mathematical Biology Program.
- HOLMES, E. E., SABO, J. L., VISCIDO, S. V., AND FAGAN, W. F. 2007. A statistical approach to quasi-extinction forecasting. *Ecology Letters* 10:1182–1198.
- HOLMES, E. E. AND WARD, E. W. 2010. Analyzing noisy, gappy, and multivariate population abundance data: modeling, estimation, and model selection in a maximum-likelihood framework. Technical report, Northwest Fisheries Science Center, Mathematical Biology Program.
- JEFFRIES, S., HUBER, H., CALAMBOKIDIS, J., AND LAAKE, J. 2003. Trends and status of harbor seals in washington state 1978-1999. *Journal of Wildlife Management* 67:208–219.
- KALMAN, R. E. 1960. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82:35–45.
- KOHN, R. AND ANSLEY, C. F. 1989. A fast algorithm for signal extraction, influence and cross-validation in state-space models. *Biometrika* 76:65–79.
- KOOPMAN, S. J. 1993. Distrubance smoother for state space models. *Biometrika* 80:117–126.
- KOOPMAN, S. J., SHEPHARD, N., AND DOORNIK, J. A. 1999. Statistical algorithms for models in state space using ssfpack 2.2. *Econometrics Journal* 2:113–166.
- LELE, S. R., DENNIS, B., AND LUTSCHER, F. 2007. Data cloning: easy maximum likelihood estimation for complex ecological models using bayesian markov chain monte carlo methods. *Ecology Letters* 10:551–563.
- MCLACHLAN, G. J. AND KRISHNAN, T. 2008. The EM algorithm and extensions. John Wiley and Sons, Inc., Hoboken, NJ, 2nd edition.
- RAUCH, H. E. 1963. Solutions to the linear smoothing problem. *IEEE Transactions on Automatic Control* 8:371–372.
- RAUCH, H. E., TUNG, F., AND STRIEBEL, C. T. 1965. Maximum likelihood estimation of linear dynamical systems. *Journal of AIAA* 3:1445–1450.
- SCHWEPPE, F. C. 1965. Evaluation of likelihood functions for Gaussian signals. *IEEE Transactions on Information Theory* IT-r:294–305.

- SHUMWAY, R. AND STOFFER, D. 2006. Time series analysis and its applications. Springer-Science+Business Media, LLC, New York, New York, 2nd edition.
- SHUMWAY, R. H. AND STOFFER, D. S. 1982. An approach to time series smoothing and forecasting using the EM algorithm. *Journal of Time Series Analysis* 3:253–264.
- STAPLES, D. F., TAPER, M. L., AND DENNIS, B. 2004. Estimating population trend and process variation for PVA in the presence of sampling error. *Ecology* 85:923–929.
- STOFFER, D. S. AND WALL, K. D. 1991. Bootstrapping state-space models: Gaussian maximum likelihood estimation and the Kalman filter. *Journal of the American Statistical Association* 86:1024–1033.
- TAPER, M. L. AND DENNIS, B. 1994. Density dependence in time series observations of natural populations: estimation and testing. *Ecological Monographs* 64:205–224.
- WARD, E. J., CHIRAKKAL, H., GONZÁLEZ-SUÁREZ, M., AURIOLES-GAMBOA, D., HOLMES, E. E., AND GERBER, L. 2009. Inferring spatial structure from time-series data: using multivariate state-space models to detect metapopulation structure of California sea lions in the Gulf of California, Mexico. *Journal of Applied Ecology* 1:47–56.
- ZUUR, A. F., FRYER, R. J., JOLLIFFE, I. T., DEKKER, R., AND BEUKEMA, J. J. 2003. Estimating common trends in multivariate time series using dynamic factor analysis. *Environmetrics* 14:665–685.



---

## Index

- animal tracking, 111
  - kftrack, 117
- bootstrap
  - innovations, 11, 28, 29
  - MARSSboot function, 11, 45
  - parametric, 11, 28, 29
- confidence intervals, 66
  - Hessian approximation, 11, 66
  - MARSSparamCIs function, 11
  - non-parametric bootstrap, 11
  - parametric bootstrap, 11, 66
- density-independent, 49
- diagnostics, 77
- DLM package, 3
- error
  - observation, 50
  - process, 49, 50
- errors
  - degenerate, 5
  - ill-conditioned, 5
- estimation, 53
  - BFGS, 32
  - Dennis method, 54
  - Kalman filter, 10, 25
  - Kalman smoother, 10, 25
  - Kalman-EM, 10, 27, 53
  - maximum-likelihood, 53, 54
  - Newton methods, 28
  - quasi-Newton, 10, 32
  - REML, 3
- extinction, 49
  - diffusion approximation, 58
  - uncertainty, 62
- functions
  - as.marssm, 124
  - checkPopWrap, 124
  - is.marssm, 11, 124
  - is.marssMLE, 10
  - MARSS, 9, 32, 34–37, 124
  - MARSSaic, 10, 28, 29, 46, 126
  - MARSSboot, 11, 28, 43, 45
  - MARSShessian, 11, 126
  - MARSSkem, 10, 27, 125, 126
  - MARSSkf, 10, 25, 26, 40
  - MARSSmcinit, 10, 28
  - MARSSoptim, 10
  - MARSSparamCIs, 4, 11, 28, 39, 126
  - MARSSsimulate, 11, 29, 43, 44
  - MARSSvectorizeparam, 11, 40, 44
  - optim, 3, 10, 28
  - popWrap, 124
  - summary, 11, 38
- initial conditions
  - setting for BFGS, 33
- KFAS package, 3
- likelihood, 10, 26, 44
  - and missing values, 27
  - innovations algorithm, 26
  - MARSSkf function, 44
  - missing value modifications, 26

- multimodal, 28
  - troubleshooting, 5, 28
- MARSS model, 1, 111
  - DFA example, 97
  - multivariate example, 69, 89, 111
  - print, 38
  - summary, 38
  - univariate example, 49
- missing values, 3
  - and AICb, 29
  - and parametric bootstrap, 28
  - likelihood correction, 27
- model selection, 29, 89
  - AIC, 29, 75, 77, 86, 87, 92, 96
  - AICc, 29, 86
  - bootstrap AIC, 29, 86
  - bootstrap AIC, AICbb, 29, 46
  - bootstrap AIC, AICbp, V, 29, 46, 86
  - MARSSaic function, 10, 46
- model specification
  - in MARSS, 13
  - in marssm objects, 21
- objects
  - marssm, 9, 11, 123
  - marssMLE, 9, 45, 125, 126
  - popWrap, 124
- simulation, 29, 43, 50
  - MARSSsimulate function, 11, 43, 44
- standard errors, 11